

# CFD Simulations of the Space Launch System Ascent Aerodynamics and Booster Separation

Stuart E. Rogers\*

*NASA Ames Research Center, Moffett Field, CA 94035, USA*

Derek J. Dalle†

*Science and Technology Corp., Moffett Field, CA 94035, USA*

William M. Chan‡

*NASA Ames Research Center, Moffett Field, CA 94035, USA*

**This paper presents details of Computational Fluid Dynamic modeling of the Space Launch System during ascent. The primary focus of the paper is the flow simulation of the vehicle during ascent using the Overflow Navier-Stokes code. Computations of 739 first-stage flight conditions covering a range of Mach numbers, angles of attack, and roll angles were computed. The overset grid system contained 375 million grid points, and over 28 million CPU hours were used in the simulations. The simulations were run on the Pleiades supercomputer at the NASA Advanced Supercomputer Center at Ames Research Center. The data products from this work include integrated line-loads, surface pressure coefficients, venting pressures, and protuberance air-loads. Detailed comparisons were made of the aerodynamic performance predicted by Overflow and the wind-tunnel derived aerodynamic database. A small number of the cases were run with two different turbulence models and with two differencing schemes. These results were used to quantify the sensitivity to the choice of the turbulence model and to the differencing scheme. The paper also introduces an effort to use the inviscid, unstructured Cartesian solver Cart3D to compute the aerodynamics during booster separation. Adaptive mesh refinement is being used to enable accurate simulations of sixteen booster-separation-motor plumes. The use of this tool is explored in preparation for building a booster-separation aerodynamic database.**

## I. Introduction

The NASA Space Launch System (SLS) will provide an entirely new heavy-lift capability that enables launch of more mass to orbit than any other launch vehicle, and makes possible crew missions beyond Earth orbit. The initial SLS configuration, shown in Fig. 1, consists of a 200-foot core stage, right and left solid-rocket boosters, an interim cryogenic propulsion stage, stage adapter, and the Orion spacecraft. This configuration has a 70-metric ton payload capability, will provide 10 percent more thrust at launch than the Saturn V rocket, and carry more than three times the payload of the space shuttle.

The SLS program is tasked with developing an all new series of launch vehicles, and although many components include legacy hardware from the Space Shuttle program, this requires developing many aerodynamic related databases. Budget limitations constrain the number of wind-tunnel tests, thus the program relies on Computational Fluid Dynamics (CFD) analysis to provide a significant amount of data. Use of CFD enables a reduction in conservatism that can be translated into higher payload to orbit.

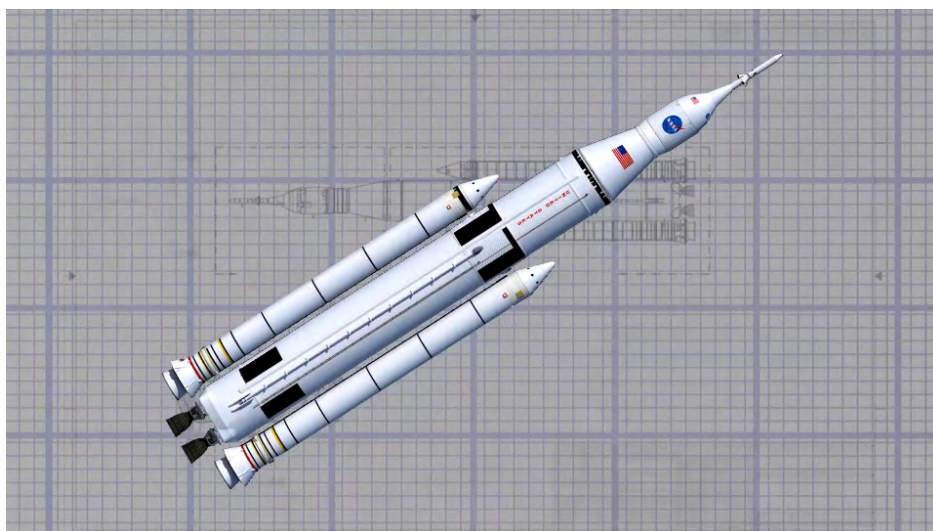
Recent CFD analysis of launch vehicle ascent aerodynamics included work for the NASA Space Shuttle program. This work included aerodynamic analysis of the events leading to the Columbia accident during the STS-107 mission [1], CFD analysis used in debris-transport analysis [2], aerodynamic analysis of proposed

---

\*Aerospace Eng., NASA Advanced Supercomputing Division, Associate Fellow AIAA, [stuart.e.rogers@nasa.gov](mailto:stuart.e.rogers@nasa.gov)

†Research Scientist/Engineer, NASA Advanced Supercomputing Division, Member AIAA, [derek.j.dalle@nasa.gov](mailto:derek.j.dalle@nasa.gov)

‡Computer Scientist, NASA Advanced Supercomputing Division, Senior Member AIAA, [william.m.chan@nasa.gov](mailto:william.m.chan@nasa.gov)



**Figure 1. Artist rendition of the SLS initial configuration**

design changes during the return-to-flight efforts, and analysis of the booster-separation plume effects [3]. The NASA Ares-I program developed aerodynamic databases using CFD analysis [4], and although this program was terminated, the Ares I-X flight test was successful. The NASA Ares-V program utilized multiple CFD solvers to produce an ascent aerodynamic database for that launch vehicle [5, 6]. The NASA Orion Multi-Purpose Crew Vehicle (MPCV) program utilized CFD analysis to develop the aerodynamic database for the launch-abort vehicle [7–11].

The SLS program is using CFD to provide data for several databases and analyses. These include vehicle sectional forces and moments used to develop the distributed line-loads database. They include the computed surface pressures over the entire vehicle to be used in venting analysis. Also included are integrated forces and moments acting on all of the individual protuberances, including feedlines and their support brackets, pressure lines and their brackets, core-to-booster attach hardware, camera mount fairings, and others. CFD analysis is also being used to develop an aerodynamic database for booster separation. The current work is all being performed for the design-analysis cycle 3 (DAC3) geometry of the vehicle. DAC3 is the configuration for Critical Design Review (CDR). This review requires that the SLS design must demonstrate appropriate maturity to support proceeding with full-scale fabrication, assembly, integration, and test.

This paper describes two of the current efforts in the SLS program. The first of these is the use of Reynolds-Averaged Navier-Stokes (RANS) CFD to simulate flow over the vehicle during ascent at hundreds of flight conditions in an effort to provide all of the data described above. The second effort is the use of an Euler CFD solver to simulate the aerodynamics during booster separation. In the following sections of the paper, the RANS computational approach is presented. The vehicle geometry preparation is described, followed by details of the grid generation. The flow-solver inputs and execution strategy are discussed. The computational resources required for these simulations are detailed, and the methods of extracting and using the data are presented. Comparisons of forces and moments from the CFD and from the wind-tunnel derived aerodynamic database are made. Some results of sensitivity studies on the choice of turbulence model and differencing scheme are given. Finally, some initial analysis of the booster separation work is discussed.

## **II. Computational Approach**

The structured overset grid approach was used in the current work because of its ability to provide high-quality viscous grids on the complicated SLS geometry, including all of the vehicle protuberances of aerodynamic interest. The NASA CFD code Overflow [12] was utilized as the flow solver. The generation of the

overset grids is time-intensive and requires significant user expertise compared to production unstructured CFD flow solvers. However, once the grids have been generated, the structured grid approach requires significantly less computational resources than unstructured solvers. Because the current work requires on the order of 1000 simulations, the overset-structured grid approach was the only approach capable of producing results in the required time frame.

The current simulations neglect the effect of the main-engine plumes and the plumes of the solid-rocket boosters. The nozzles were sealed at their exit plane, and the seals were treated as no-slip walls. Therefore the current surface pressure on the aft end of the vehicle does not include the effects of plumes, or of any plume-induced flow separation, which only occurs at high altitude just prior to booster separation.

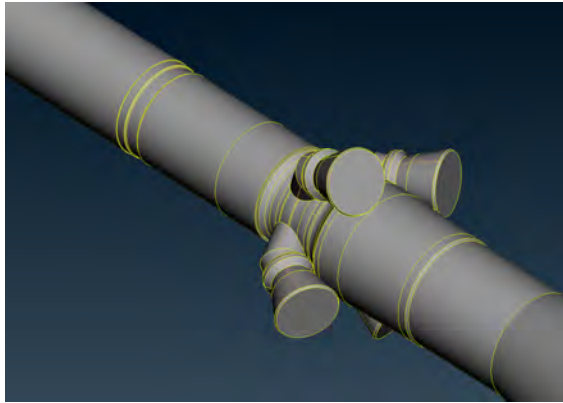
### III. Vehicle Geometry

The SLS-10005 outer mold line (OML) geometry served as the basis for the current CFD model. This model was simplified further from the original to facilitate the grid-generation process. This was necessary to keep the total computational requirements within the computational resources available and within the schedule requirements of the program. The ANSA software was used to manipulate and simplify the geometry, and then to export the geometry as very fine triangulated surfaces to be used in the surface grid-generation process. The primary components of the geometry include the Launch Abort Vehicle (LAV), the MPCV and its service module, the Launch Vehicle Stage Adapter (LVSA), the Core stage, and the two Solid-Rocket Boosters (SRB). Images of the OML geometry details are shown in the following figures. Figure 2a shows the geometry near the LAV abort-motor nozzles, and Fig. 2b shows the geometry of the MPCV boost-protection cover. Figure 2c shows the booster nosecone and forward booster separation motor nozzles. Figure 2d shows the geometry near the forward end of the LO2 feedline and the pressure lines on the core stage, whereas Fig. 2e shows the aft ends of these lines. Figure 2f illustrates the geometry on the aft end of one of the boosters.

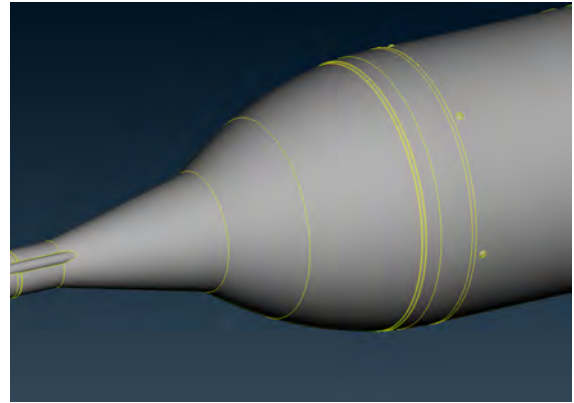
### IV. Grid Generation

The creation of the grid system for the SLS-10005 DAC3 vehicle was a team effort consisting of two experienced and two novice users of the grid tools. Taking full advantage of the overset grid approach, grids for the component parts of the complex geometry were generated in parallel by the team members. The axisymmetric main component grids (core and boosters) were created first, followed by grids for all the protuberances that reside on the main components. Utilizing macros from the Chimera Grid Tools (CGT) [13, 14] version 2.1p+ script library, scripts were created for all pre-processing steps needed prior to running the flow solver. These include the generation of surface and volume grids, creation of hole-cutting instructions for overset domain connectivity, prescription of boundary conditions in the flow solver input file, and specifications of components for forces and moments integration. The scripting approach provides a record of the all steps used in the complex grid generation process, and allows rapid automated re-generation of the entire grid system. Key inputs in the grid generation steps such as grid spacings, marching distances and stretching ratios are parameterized in the script. This enables quick grid resolution adjustments to achieve high grid quality.

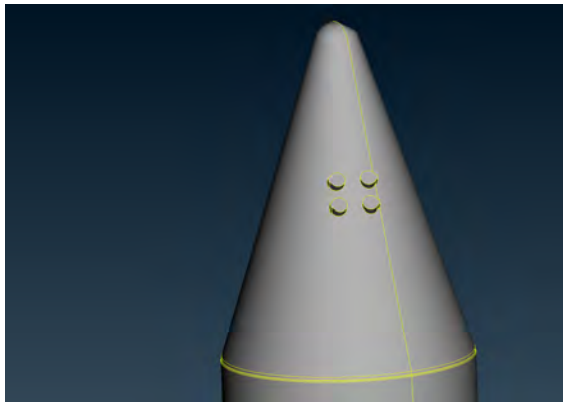
Surface grids were created using a combination of algebraic and hyperbolic methods via the *surgrd* module in CGT. A maximum surface spacing of ten inches and maximum circumferential resolution of one degree were utilized in the axisymmetric components. In the vicinity of the attach hardware and protuberances, the grids were made to be much finer, and all features were resolved with grid spacings of fractions of an inch or less. The body-fitted volume grids were generated hyperbolically using the *hypgen* module in CGT with a normal viscous wall-spacing targeted to produce the first cell within a  $Y^+$  value of 1.0 at maximum-dynamic pressure conditions. Off-body Cartesian grids were automatically generated using the multi-layer feature in the Overflow flow solver. After combining the individual body-fitted and off-body grids, the overset hole-cutting and interpolation were performed using the X-ray capability within CGT and Overflow. The completed grid system was composed of 888 body-fitted zones and 304 Cartesian off-body



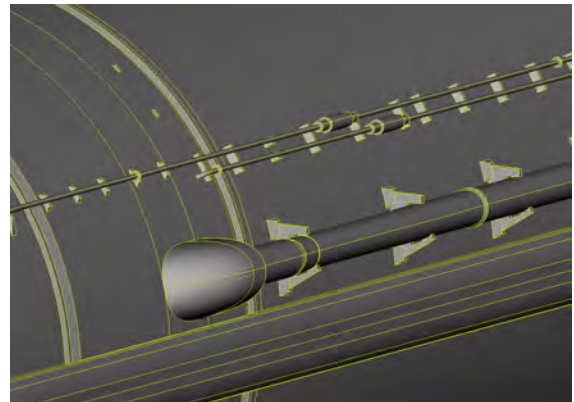
**a) LAV abort-motor nozzles**



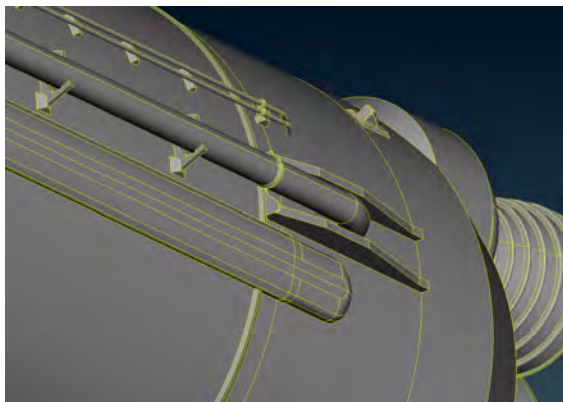
**b) MPCV boost-protection cover**



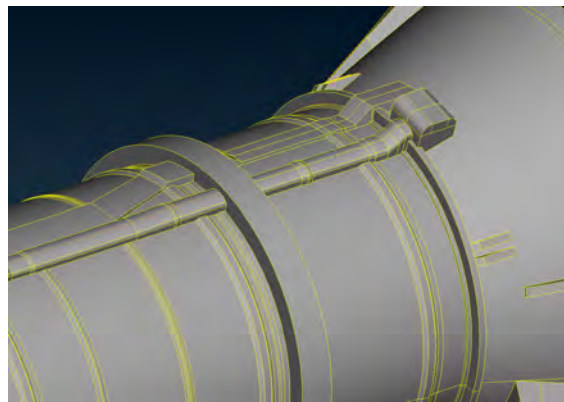
**c) Forward booster-separation motors**



**d) Forward end of feedline and pressure lines**



**e) Aft end of feedline and pressure lines**



**f) Booster aft end**

**Figure 2. Details of the SLS vehicle geometry.**

grids. The total number of grid points was just over 375 million with only 203 orphan points remaining after overset connectivity. The entire grid system was created in about two months by the four person team. It is one of the most complex and efficiently generated grid systems using overset structured grids ever produced. Figure 3a shows a side view of the surface and volume grids on the entire vehicle, Fig. 3b shows the surface grids on the forward part of the vehicle. Figure 3c shows the surface grids near the nose of the right booster, and Fig. 3d shows a view of the aft end of the right booster. Figure 3e shows the surface grids on +Z LO2 feedline and brackets, and Fig. 3f shows the surface grids on the forward-attach hardware of the left booster.

## V. Flow Conditions

The freestream conditions for Mach number ( $M$ ), total angle of attack ( $\alpha_t$ ), and velocity roll angle ( $\phi$ ) used in the simulations are:

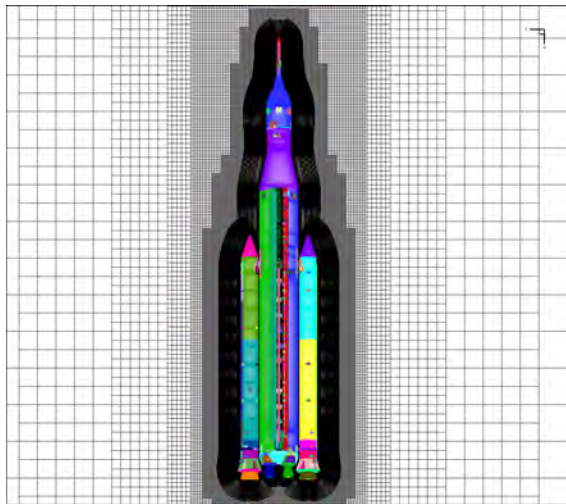
- $M = 0.7, 0.8, 0.9, 0.95, 1.05, 1.10, 1.20, 1.30, 1.40, 1.60, 1.75, 2.0$ :
  - $\alpha_t = 0.0$ ;  $\phi = 0$
  - $\alpha_t = 2.0, 4.0$ ;  $\phi = 0, 30, 45, 60, 90, 120, 135, 150, 180$
  - $\alpha_t = 2.0, 4.0$ ;  $\phi = 210, 225, 240, 270, 300, 315, 330$
- $M = 0.5, 2.5, 3.0, 3.5, 4.0, 4.5, 5.0$ :
  - $\alpha_t = 0.0$ ;  $\phi = 0$
  - $\alpha_t = 2.0, 4.0, 6.0$ ;  $\phi = 30, 45, 60, 90, 120, 135, 150, 180$
  - $\alpha_t = 2.0, 4.0, 6.0$ ;  $\phi = 210, 225, 240, 270, 300, 315, 330$

## VI. Overflow Inputs

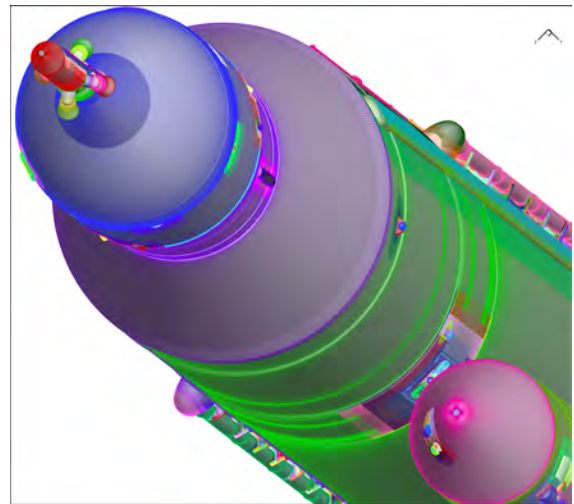
The flowfields were computed with the standard release version 2.2h of the Overflow code, using the Message-Passing Interface (MPI) parallel version. The computations were performed on the SGI Altix system known as *Pleiades* at the NASA Advanced Supercomputing (NAS) facility. The computations were all performed on the Ivy Bridge nodes of this machine. Overflow was compiled using the Intel Fortran compiler *ifort* version 2013.5.192 and the SGI MPI implementation mpt.2.10r6. All of the runs used 41 Ivy Bridge nodes, consisting of 820 cores.

Central differencing plus scalar smoothing was used, in combination with the implicit block tridiagonal approximate-factorization scheme. All viscous terms, including cross terms, were enabled in the code. The Spalart-Allmaras turbulence model [15] was used, and the flow was assumed to be turbulent everywhere. These were the baseline Overflow inputs which were used for all of the production cases in the current work. In addition, a small number of cases were run to quantify the sensitivity of the solutions to the choice of turbulence model and choice of differencing scheme. The sensitivity to the choice of turbulence model was studied by repeating some cases using the SST turbulence model [16, 17]. The same set of cases was also run using a third-order HLLC upwind-differencing scheme, together with an SSOR line-relaxation implicit scheme.

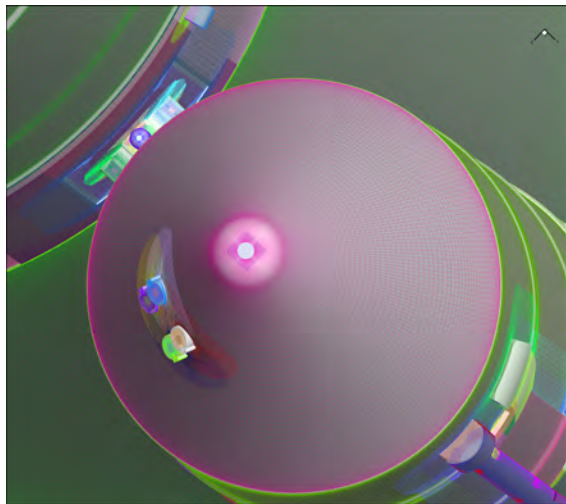
For the initial start-up phase and for steady-state iterations, first-order time integration was used together with locally-scaled time-step sizes based on a constant Courant-Friedrichs-Lewy (CFL) number. The time-accurate integration used a dual time-stepping approach in which the outer-level integration was second-order in time. The inner-level integration used ten sub-iterations per time step with a non-dimensional time-step size of 1.0. This corresponds to a physical time-step size of  $1.65 \times 10^{-5}$  to  $1.5 \times 10^{-4}$  seconds for free-stream Mach numbers ranging from 5.0 to 0.5, respectively. These values were determined after running a small number of test cases varying both the time-step size and the number of inner sub-iterations. This combination provided reasonably converged results relative to the cost of running more inner sub-iterations and/or running with a smaller time-step size.



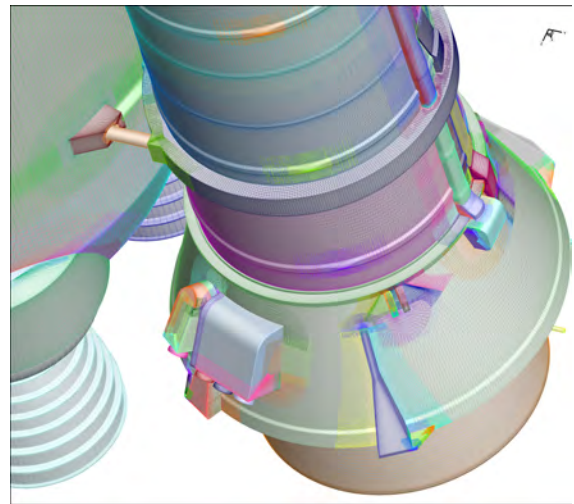
**a) Side view of entire grid system**



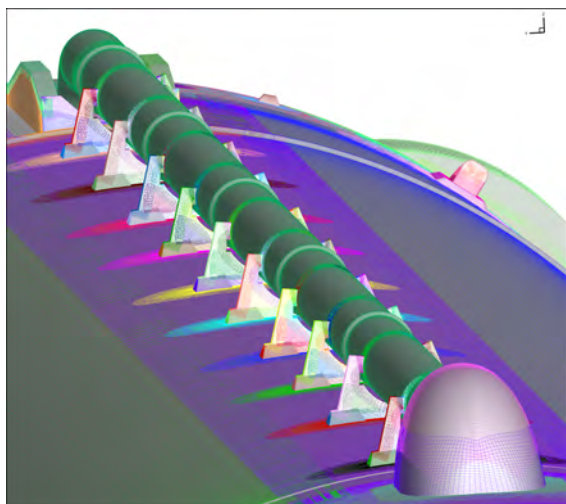
**b) Forward vehicle surface**



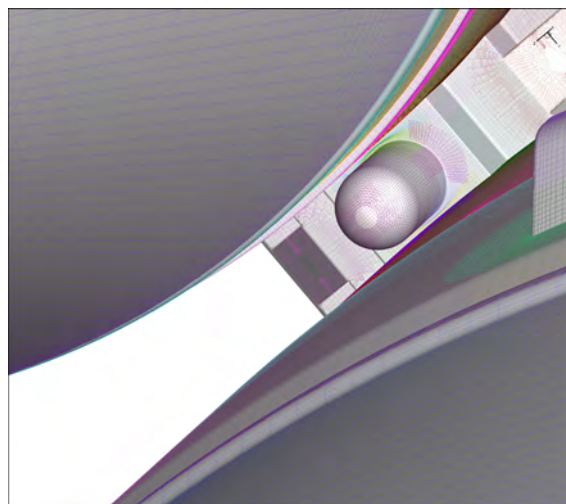
**c) SRB nosecap**



**d) Booster aft end**



**e) Liquid oxygen feedline on the core**



**f) Forward attach hardware**

**Figure 3. Different views of the surface grids on portions of the SLS vehicle.**

## VII. Overflow Execution and Convergence

The following run sequence was found to be the most efficient strategy for the SLS DAC3 calculations. The code was run with a sequence of different inputs, starting with a full-multi-grid sequencing for the first 4000 steady-state iterations, using relatively higher dissipation and smaller time-step sizes. Over the next 6000 steady-state iterations, the time-step size was increased and the dissipation coefficients were decreased in three increments. This was followed by 6000 to 10,000 additional steady-state iterations. If the case converged to a steady-state, the simulation was terminated. The case was considered converged if the L2-norm of the residual had converged at least three to four orders of magnitude, and the aerodynamic force coefficients converged to a steady state. Several different force coefficients were monitored, including the forces acting on the total core, each SRB, large protuberances like the liquid-oxygen feedline, and smaller protuberances like the aft booster-separation motors on the SRBs.

If the steady-state iterations did not converge or the forces were unsteady, the case was continued with time-accurate integration. The time-accurate integration was run until the time-averaged mean of the monitored force coefficients were converged. Then the final solution was extracted from a time average of the dependent variables over a span of at least 2000 time steps. During the Overflow runs, the code computed the forces and moments on all components of the vehicle every ten iterations to enable the convergence monitoring. The moments were all computed about a reference center at the base of the vehicle.

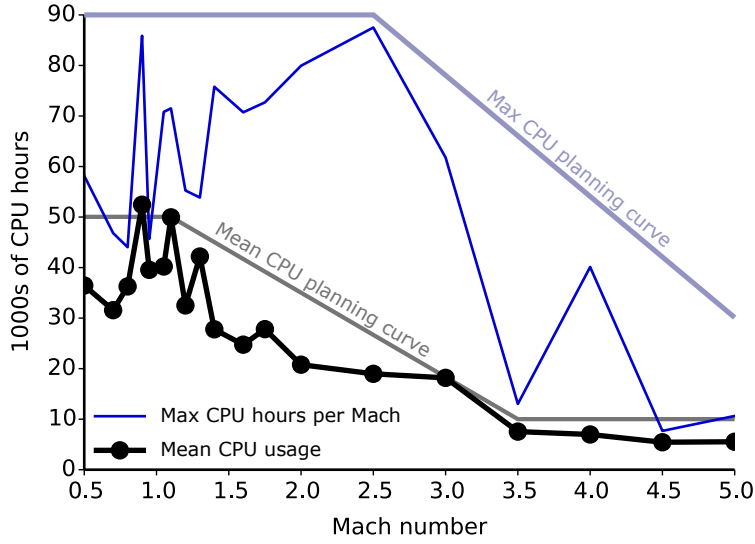
The NASA-developed *overlst* software was used to execute, manage, and post-process all of the Overflow runs. This is a script written in the Perl language. The Overflow community has yet to devise a reliable automatic method for determining convergence, and so the key to implementing this run strategy for hundreds of Overflow simulations is the ability to easily monitor the progress of each case. The monitoring was greatly aided by the use of the *overlst* software, which greatly simplifies the creation of batches of convergence plots which can be easily and rapidly examined by the user. This software also provides a simple interface for the user to control batches of simulations, adding more iterations/time-steps, and switching from steady-state to time-accurate integration.

### A. Computational Resources

Figure 4 shows a summary of the CPU hours required to run the ascent database simulations, with the cases grouped by Mach number. At each Mach number there are either 14 (from Mach 0.7 to 2.0) or 21 (otherwise)  $\alpha_T, \phi$  combinations in the run matrix. The black curve with dots in Fig. 4 shows the average CPU-hour requirement per case at the given Mach number, and the blue curve is the maximum number of CPU-hours per case at that Mach number.

Several things stand out in Fig. 4. Most importantly, the cases around Mach 1.0 require the most CPU hours, which is expected since this is the most likely regime for unsteady flow. Second, the CPU requirements are elevated (compared to those of Mach 5.0) even at true supersonic conditions as high as Mach 3.0. This is because there is not a clean break in Mach number where all lower-Mach cases require time-accurate solution and all higher-Mach cases are steady. For low supersonic cases, there are often orientations that have either vortex shedding, unsteady shock interactions, or regions of unsteady flow separation. In fact, the single case requiring the most CPU hours was at Mach 2.5, despite the fact that the average usage at Mach 2.5 was quite low. Finally, it is apparent that applying the most conservative run sequence (which required about 85,000 CPU hours) to all cases would dramatically increase the total requirements for the entire run matrix. The total CPU usage for this subset of the run matrix was 3.2 million CPU hours, whereas the number would have been over 27 million if each case had required 85,000 CPU hours.

The simple piecewise linear curves labeled “planning curves” in Fig. 4 are a simplified interpretation of the data. It was observed that a transonic case takes, on average, about five times as much CPU time as a high supersonic case. Similarly, it was estimated that a worst-case Mach 5.0 case could take as much as three times more CPU hours than a nominal case, and a worst-case transonic case could take as much as 9 times as long as the nominal Mach 5.0 case. If  $N_{M5}$  is the nominal number of CPU hours per case at Mach



**Figure 4. CPU usage for ascent database cases grouped by Mach number**

5 (in this case 10,000 CPU hours), then the two curves are

$$N_{(mean,planning)} = \begin{cases} 5N_{M5} & : M \leq 1.1 \\ (5 - \frac{5}{3}(M - 1.1))N_{M5} & : 1.1 < M \leq 3.5 \\ N_{M5} & : M > 3.5 \end{cases} \quad (1)$$

$$N_{(max,planning)} = \begin{cases} 9N_{M5} & : M \leq 2.5 \\ (9 - 2.4(M - 2.5))N_{M5} & : 2.5 < M \leq 5.0 \\ 3N_{M5} & : M > 5.0 \end{cases} \quad (2)$$

These functions are expected to be valid for most launch vehicle ascent simulations using Overflow, although geometry (for example, the presence or absence of large protuberances) can affect the prevalence of unsteady cases, and significant changes to Overflow inputs will invalidate this calibration. The approach (with additional calibrations to the slopes of the lines, etc.) may also be useful for other types of problems and other flow solvers, but additional considerations may be relevant. In addition to checking intuition, these functions can be useful for planning, which is discussed in the following subsection.

## B. Computational Planning Insights

Typically an estimate of computational resources for a large database such as this one are based on Rough Order of Magnitude (ROM) analyses that are thought to be conservative. Especially for vehicles or run matrices outside the experience of the organization completing the analysis, constructing a more accurate Basis Of Estimate (BOE) is difficult or impossible. For example, using a conservative estimate of the requirements at Mach 5.0 may lead to an underestimate of the total requirements for the matrix.

Based on the actual requirements for 315 cases as reported in Fig. 4, a relatively simple planning method has been developed where the CPU requirements of a case are a function of the freestream Mach number. First estimate the expected CPU requirements for the nominal inputs, for example by running a single case through a set of inputs that is sufficient to converge most high-supersonic cases or using a ROM of the CPU requirements for a high-supersonic case based on previous experience. Using a high-supersonic case such as Mach 5.0 is useful to anchor the estimate even if it is not in the run matrix because the convergence behavior in this regime is more predictable.

The second step is to generate the “Mean CPU planning curve” as in Fig. 4 based on the estimate for CPU hours at Mach 5.0. This simple function then gives an estimate of the average number of CPU hours per case as a function of Mach number,  $f_{BOE}(M)$ , and combining this information with the planned run matrix leads to an estimate of the total required number of CPU hours.

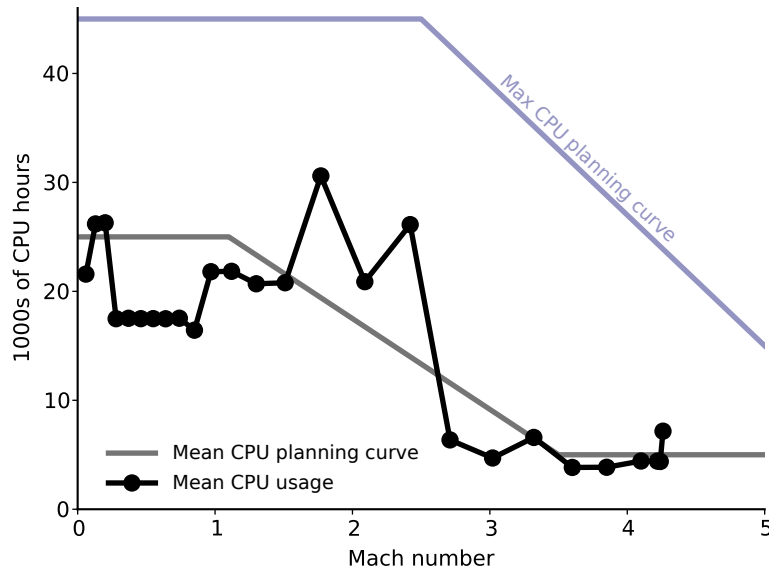


Figure 5. CPU usage for coarser grid system used for debris transport analysis

Figure 5 shows an example of this procedure applied to the coarser grid system used for debris transport analysis. Since this system has roughly half as many grid points as the full ascent grid system, a reasonable estimate of the CPU requirements for a Mach 5.0 case is 5000 CPU hours per case (half of the 10,000 for the full ascent grid system). This is all the information needed to create the “Mean CPU planning curve” and “Max CPU planning curve” in Fig. 5. The run matrix for the debris transport analysis was one condition at each Mach number marked with a dot in Fig. 5 except for Mach 3.32, at which 13 cases were run. Then the total estimate of CPU hours required would be

$$CPU_{BOE} = f_{BOE}(0.06) + f_{BOE}(0.13) + \dots + f_{BOE}(3.02) + 13f_{BOE}(3.32) + \dots + f_{BOE}(4.26) \quad (3)$$

which gives an estimate of 509,000 CPU hours. The actual CPU usage was 483,000 CPU hours, so at least in this case, the method gave an accurate prediction.

The “Max CPU planning curve” can also be useful in several cases. For one, if the run matrix has only a small number of conditions, using an average is more likely to underpredict actual CPU requirements, and using the maximum is a more appropriate conservative estimate. Second, if the computing environment is such that all or nearly all cases in the run matrix can be executed simultaneously, the worst case in the entire run matrix provides the best estimate of the wall time required.

When it is not possible to run all or nearly all cases simultaneously, the recommended procedure is to start with transonic and subsonic cases so that any outliers requiring more time to converge will have a chance to do so while running the supersonic cases. This leads to a lower probability of having one or two cases that are still running hours or days after the rest of the run matrix has been completed.

## VIII. Data Extraction and Delivery

### A. Sectional Loads Data

A key deliverable from the aerodynamics team for structural analysis is the integrated loads on slices of the vehicle. That is, the vehicle is first sliced by a number of equally-spaced axial planes, and the forces

and moments on these slices are integrated. The most important components of these loads are the normal and lateral forces, which are the most likely to bend the core stage and two boosters. To calculate these aerodynamic loads, the following steps are taken.

- The wall and wall-adjacent layers of the solution file are extracted; the second layer is needed to compute the viscous forces.
- Four sets of cuts are made: one for the core, one for each booster, and one for the entire stack. Each of these four is cut into 201 slices using the `triload` program [18] from the CGT package.
- The `triload` program constructs surface triangles covering each slice.
- Forces and moments are calculated for each slice by integrating the pressure- and viscous-force contributions at each triangle.

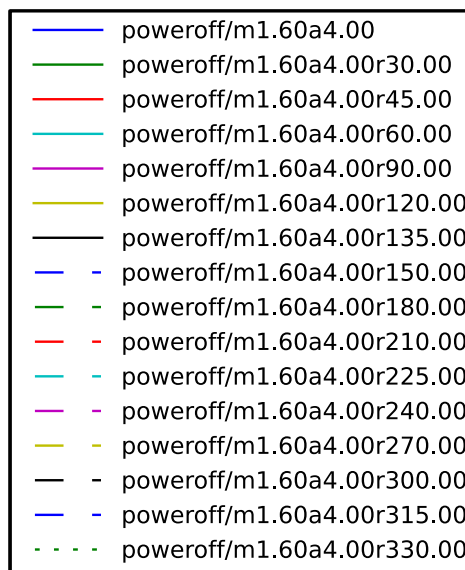
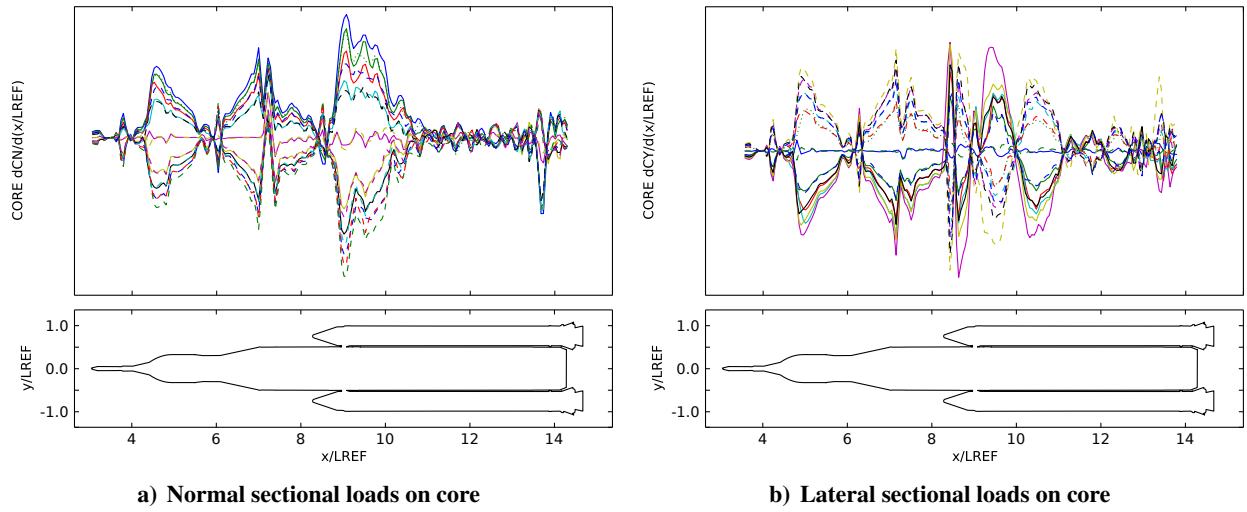


Figure 6. Legend for plots of sectional loads

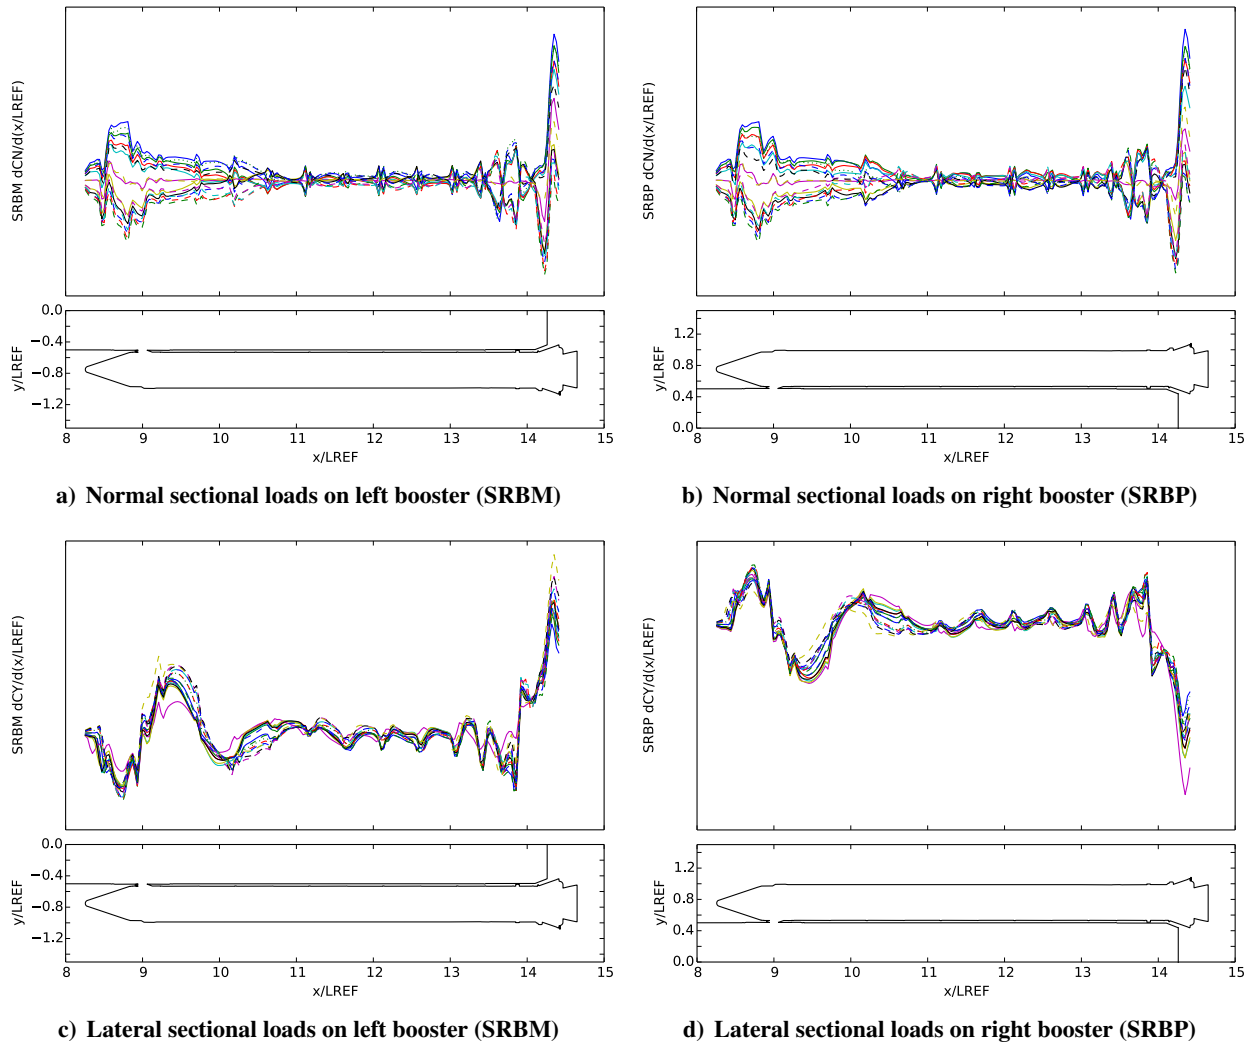
Figure 6 is a legend for the example sectional line loads plots shown in Figs. 7 and 8. These figures plot the sectional loads as a functional of the axial position on the vehicle, scaled by the reference length LREF. The subset of the run matrix used for these examples is a roll sweep at Mach 1.60 with a fixed total angle of attack at  $4^\circ$ . The sectional loads in Fig. 7 show the expected symmetry: a sign-change in sideslip angle has little effect on normal forces, and a sign-change in angle of attack (not total angle of attack) has little effect on lateral forces. In Fig. 8, the plots are organized such that plots in the same row should be either symmetric or anti-symmetric. The normal loads on the left booster (SRBM) at conditions of  $(\alpha, \beta)$  should be similar to the normal loads on the right booster (SRBP) at  $(\alpha, -\beta)$ . Unlike the core, there is no symmetry in  $\beta$  because the presence of the core on only one side of the booster means that the sign of  $\beta$  plays a bigger role. Similarly, the lateral loads on SRBM at  $(\alpha, \beta)$  are close to the opposite of the lateral loads on SRBP at  $(\alpha, -\beta)$ .

## B. Protuberance Data

Force data on each of the 140 separate protuberances was extracted. These protuberances are essentially anything other than the main booster and core bodies; they include parts like attach hardware, cameras,



**Figure 7. Sectional aerodynamic loads on the core at Mach 1.6 and 4° total angle of attack**



**Figure 8. Sectional aerodynamic loads on the solid rocket boosters at Mach 1.6 and 4° total angle of attack**

system tunnels, and pressure lines. Extracting the forces on these protuberances is the first step in a process that ensures that the protuberances remain safely attached and are able to withstand the ascent environment.

The data that is delivered to the protuberance team takes the form of triangulated surface data on each protuberance. To obtain this, The `mixsur` program from CGT is used to create a unique triangulation to cover the surface of each protuberance. Each protuberance is further broken down into individual faces, such as the top, the bottom, and the sides. A program called `overint` from the CGT package is used to extract the surface pressure coefficient ( $C_p$ ) and force components independently. The final post-processing step is the conversion of this data into two Tecplot files. One of these files contains the  $C_p$ . The other file contains the three components of the dimensionalized force (broken down into the pressure, viscous, momentum, and total contributions) and area on each triangle.

Figure 9 shows an example of axial force extraction on a forward-facing camera on the aft portion of the core, which is visible in Fig. 2e. The quantity that determines the color is the axial component of force on each triangle, and thus smaller triangles have less force.

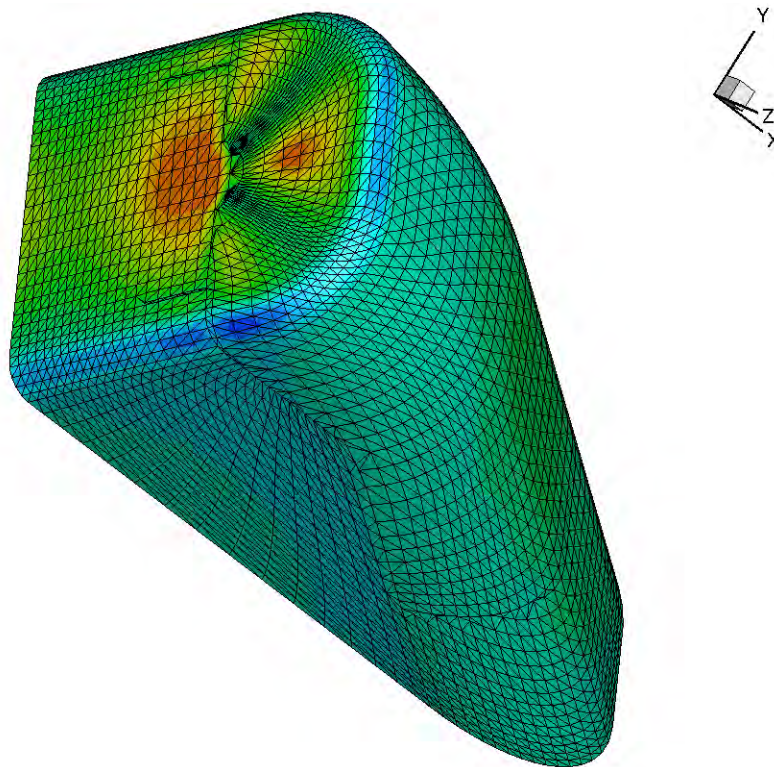


Figure 9. Axial force on each triangle of the core aft camera 1 at Mach 1.2 with  $\alpha_t = 2^\circ$  and  $\phi = 45^\circ$

### C. Surface Pressure Data

The surface  $C_p$  data on the entire vehicle was extracted and saved as formatted Tecplot data files. Each data file contains 61 different Tecplot zones corresponding to different components and surfaces on the vehicle. A perl script was used to create the  $C_p$  data file for each case. The script uses the `overint` and the `triged` programs from the CGT package to create a surface triangulation together with the surface  $C_p$  data.

### D. Boundary Layer Profiles

Boundary layer profiles were extracted to support two independent tasks: venting analysis and scaling buffet forcing functions from wind tunnel data. To extract these boundary layer profiles, the `vpro` program from

the CGT package was used to extract all state variables along a line segment perpendicular to the surface at about 40 locations on the surface. The points at which the boundary layer profiles are extracted are the physical locations of 33 vents and the locations of several transducers in an unsteady aerodynamics wind tunnel test. The data in these boundary layer profiles includes pressure, density, and the three components of velocity. The axis for the  $x$ -component of the velocity is the projection of the vehicle  $x$ -axis into the local tangent plane, the  $z$ -component is normal to the surface, and the  $y$ -component completes the right-handed system. Figure 10 shows the dimensional boundary profiles of the velocity magnitude at one vent location for all cases in the run matrix that have zero total angle of attack. there is a tendency for the boundary layer to grow as the freestream pressure decreases and the Reynolds number decreases. The slight kinks seen at the outer edges of the velocity data is due to the transition in grid resolution as the profiles cross overlapping grid boundaries.

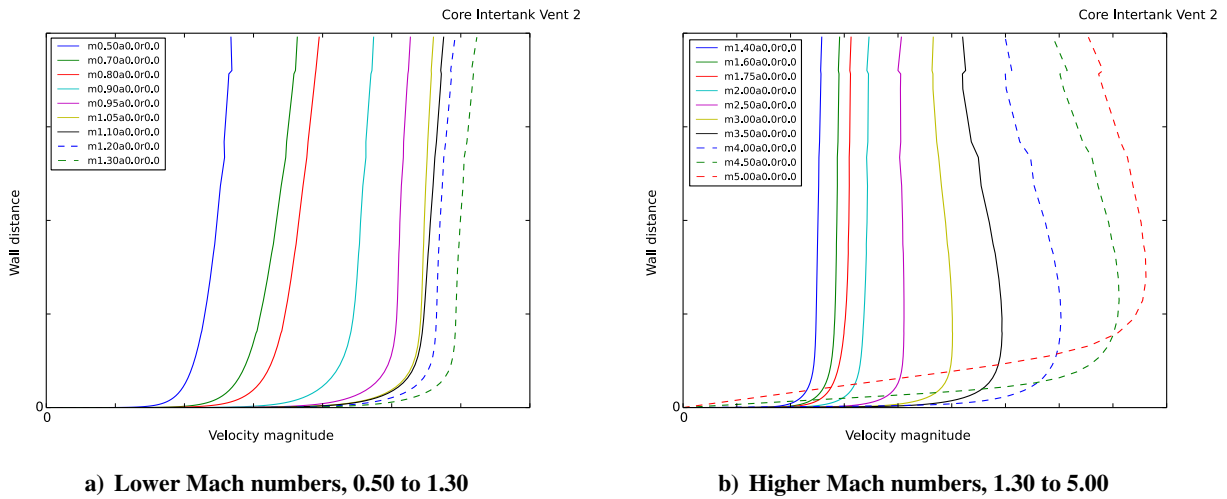


Figure 10. Boundary profiles at a vent on the core stage for Mach sweep of conditions at  $0^\circ$  total angle of attack

## IX. Aerodynamic Data

In this section the integrated force coefficients from the current CFD runs are compared to the aerodynamic databook values for the vehicle. The databook values are all based on experimental data, and include uncertainty increments to account for experimental uncertainty and repeatability, and differences between the wind-tunnel results and the flight vehicle. The Overflow CFD data in the following figures excludes the base region of the core to best match the experimental model, whose base was part of the sting mount and was non-metric. It also excludes the base and nozzle of the boosters aft of the skirt. The experiment did not include plumes or any type of plume simulator, so the aerodynamic database is designed to exclude the plume effects. Note that the database excluded the experimental data for Mach = 1.05, 1.1, and 1.2 due to shock reflection issues, and that data in this Mach range was interpolated from surrounding data.

Figures 11 through 16 show plots of the normal-force coefficient ( $C_N$ ) versus Mach number at  $\phi = 0, 30, 45, 60, 300, 315, 330$  degrees, covering the range of positive angle of attack ( $\alpha$ ). Note that while the vehicle is not symmetric about the  $Z=0$  plane, nearly identical mirrored results are seen for the other values of  $\phi$ . In these and subsequent plots the data is illustrated for  $\alpha_t = 0, 2$ , and 4 degrees. The shaded bands represent the aerodynamic database, where the black line in the center is the nominal value, and the shaded region represents the uncertainty range on either side of the nominal. The red, green, and blue colored lines with circular symbols are the time-averaged CFD values, and the surrounding color lines without symbols are the minimum and maximum values of the CFD data within the time-averaging window. It can be seen that all of the CFD time-averaged values fall within the database uncertainty range. For Mach numbers above 2.0

the CFD data agrees very well with the nominal database values. The biggest discrepancies occur for  $\alpha_t = 4$  degrees in the low supersonic range, where the CFD predicts a decrease in CN below the transonic values.

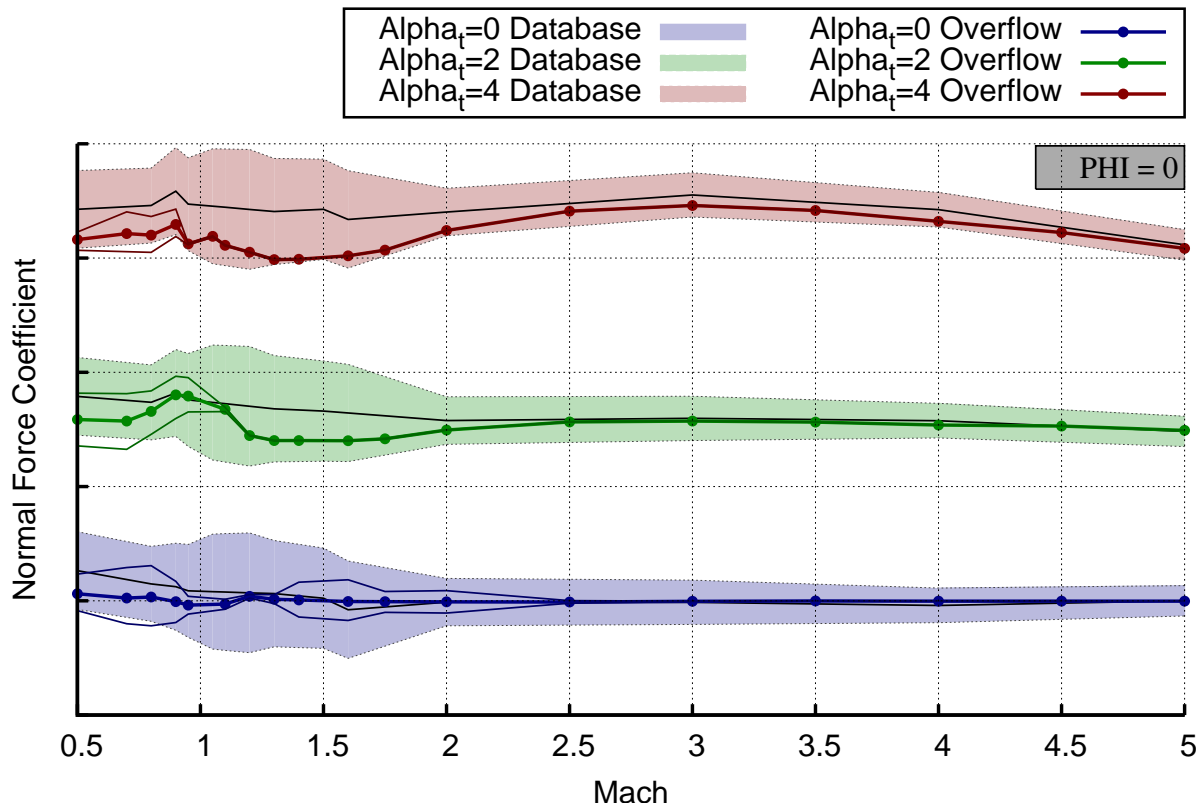


Figure 11. SLS aerodynamic database comparison of CN versus Mach number for  $\phi = 0$

Figures 17, 18, and 19 contain the plots of side-force coefficient (CY) versus Mach number for  $\phi = 45$ , 60, and 90 degrees. Very good agreement is seen between the CFD time-averaged results and the nominal database. The magnitude of the unsteady oscillations in the CFD is seen to be relatively small.

Figures 20, 21, and 22 contain the plots of pitching-moment coefficient (CLM) versus Mach number for  $\phi = 0$ , 45, and 180 degrees. Very good agreement is seen between the CFD time-averaged results and the nominal database for Mach numbers of 2.0 and above. Similar to the CN results, the CFD predicts a significant gradient through the transonic range that is not seen in the database. The magnitude of the unsteady oscillations in the CFD is seen to be relatively small, and is only non-zero at the subsonic and transonic Mach numbers.

Figures 23 through 25 are plots of the yawing-moment coefficient (CLN) versus Mach for  $\phi = 45$ , 60, and 90 degrees. The agreement between the CFD data and the nominal database values is excellent. Although not shown here, the agreement between the CFD CLN and the database CLN is excellent for all  $\phi$  angles.

Figure 26 is a plot of the rolling-moment coefficient (CLL) versus Mach number for  $\phi = 45$  degrees. It can be seen that the unsteady oscillations in the CFD data are relatively large compared to the database uncertainty. While the CFD values fall within the database uncertainty, the agreement at  $\alpha_t = 4.0$  degrees is not very good.

The differences between the CFD data and the nominal aerodynamic data can be due to a number of factors. These factors include limitations in the accuracy of the CFD, including inherent assumptions in unsteady RANS modeling and the ability of the turbulence model to reproduce the flow physics. These also include the large differences in Reynolds number between the wind-tunnel test and the flight vehicle, the possibility of strong wind-tunnel wall effects in the transonic and low supersonic ranges, and sting effects. The geometry of the wind-tunnel model is of relatively low fidelity due to the 0.8% scale of the model, and

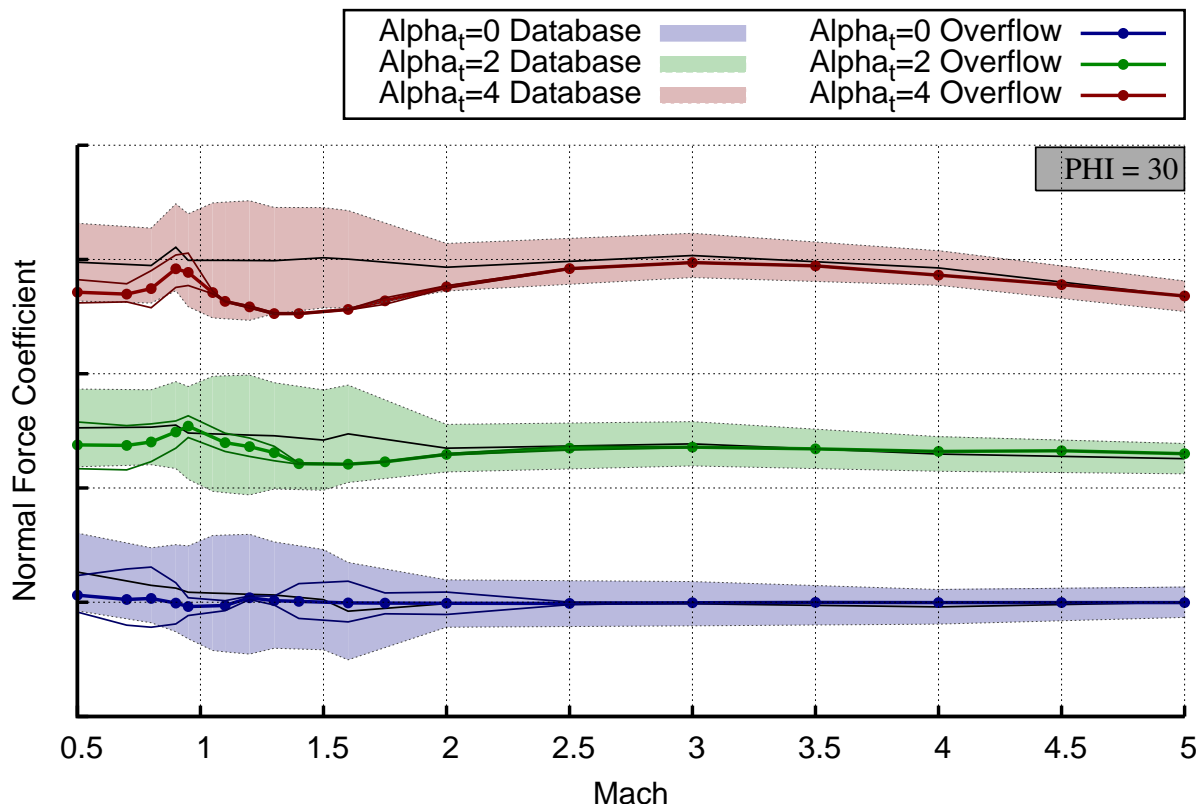


Figure 12. SLS aerodynamic database comparison of CN versus Mach number for  $\phi = 30$

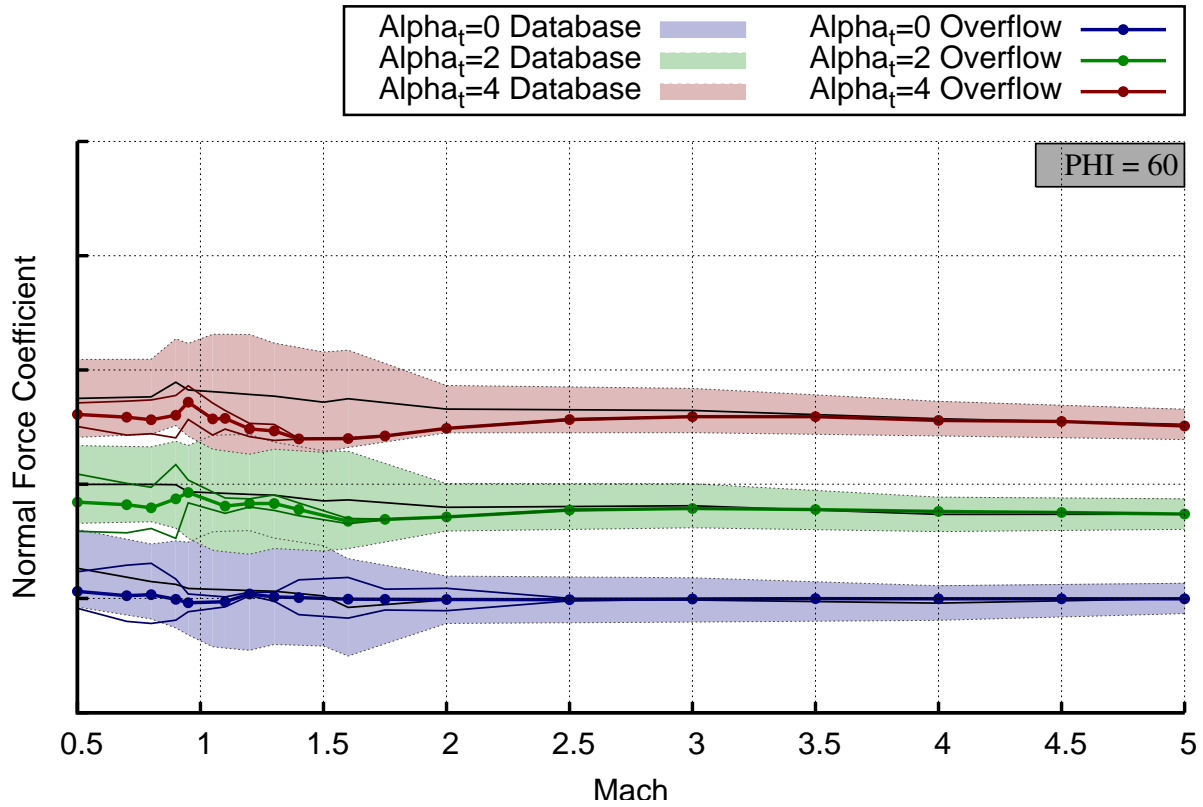


Figure 13. SLS aerodynamic database comparison of CN versus Mach number for  $\phi = 60$

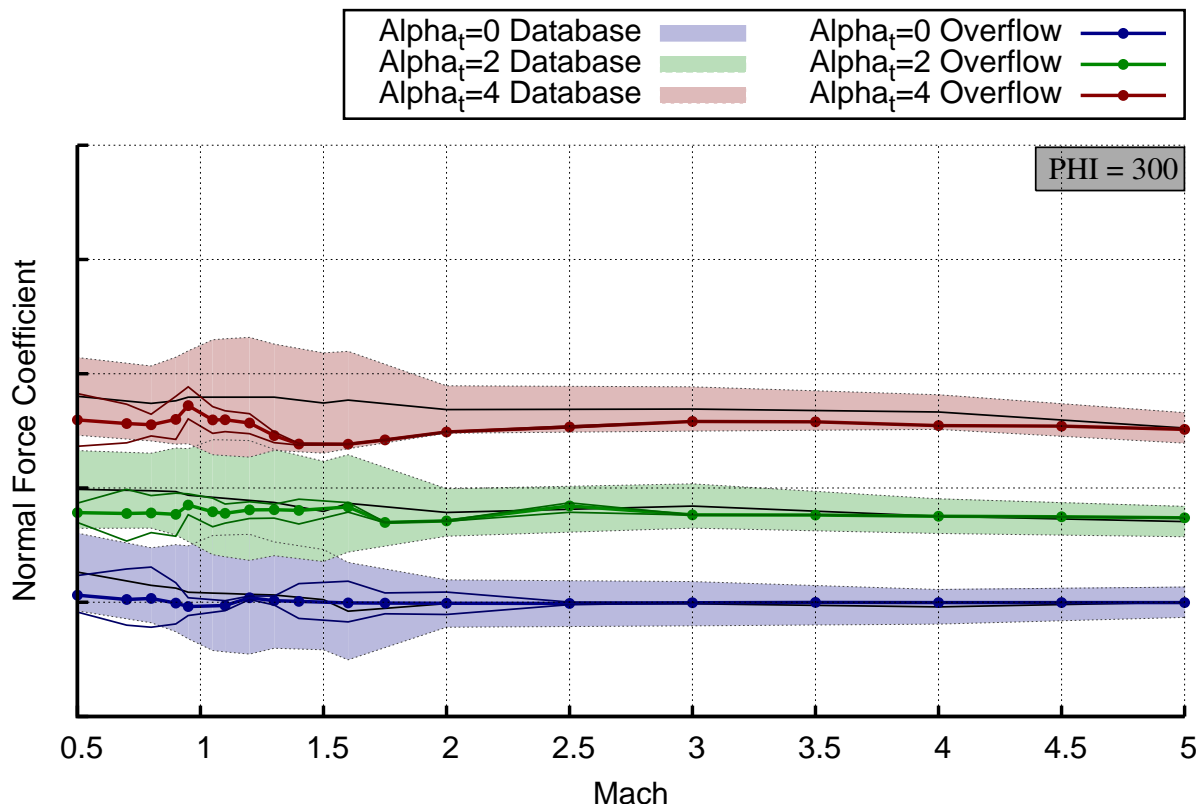


Figure 14. SLS aerodynamic database comparison of CN versus Mach number for  $\phi = 300$

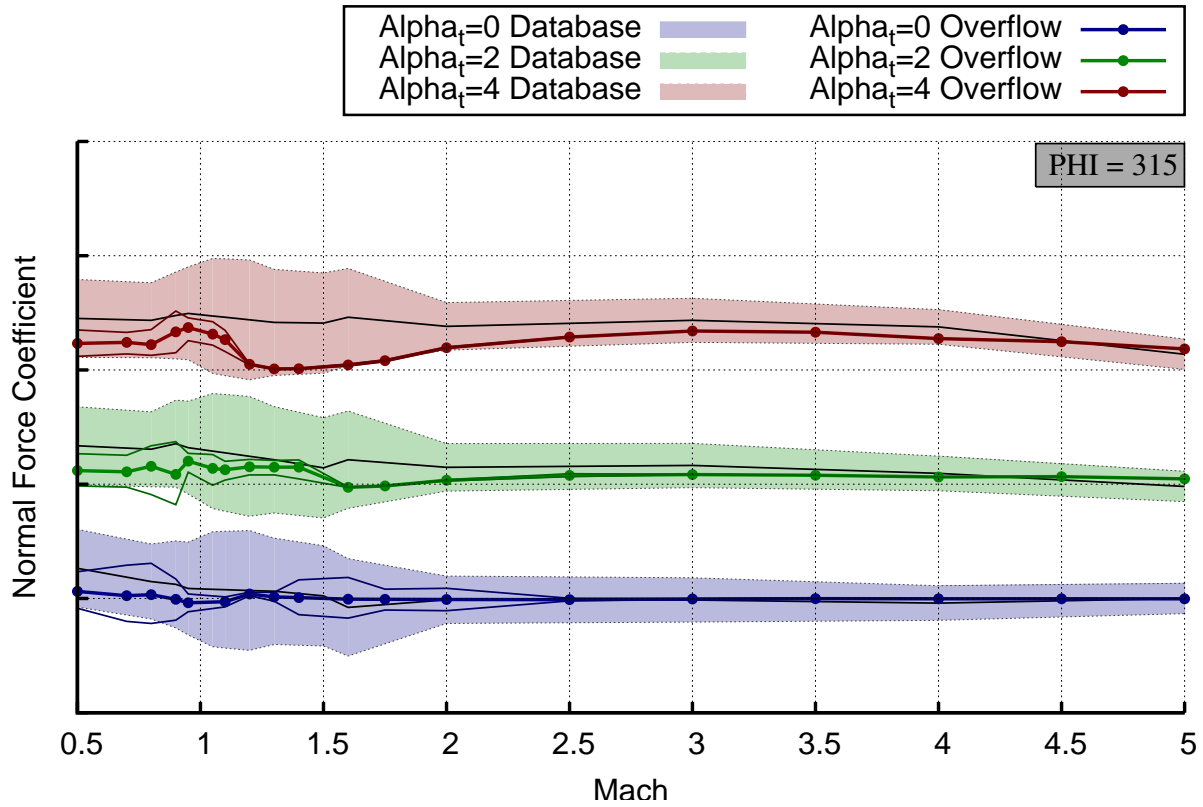


Figure 15. SLS aerodynamic database comparison of CN versus Mach number for  $\phi = 315$

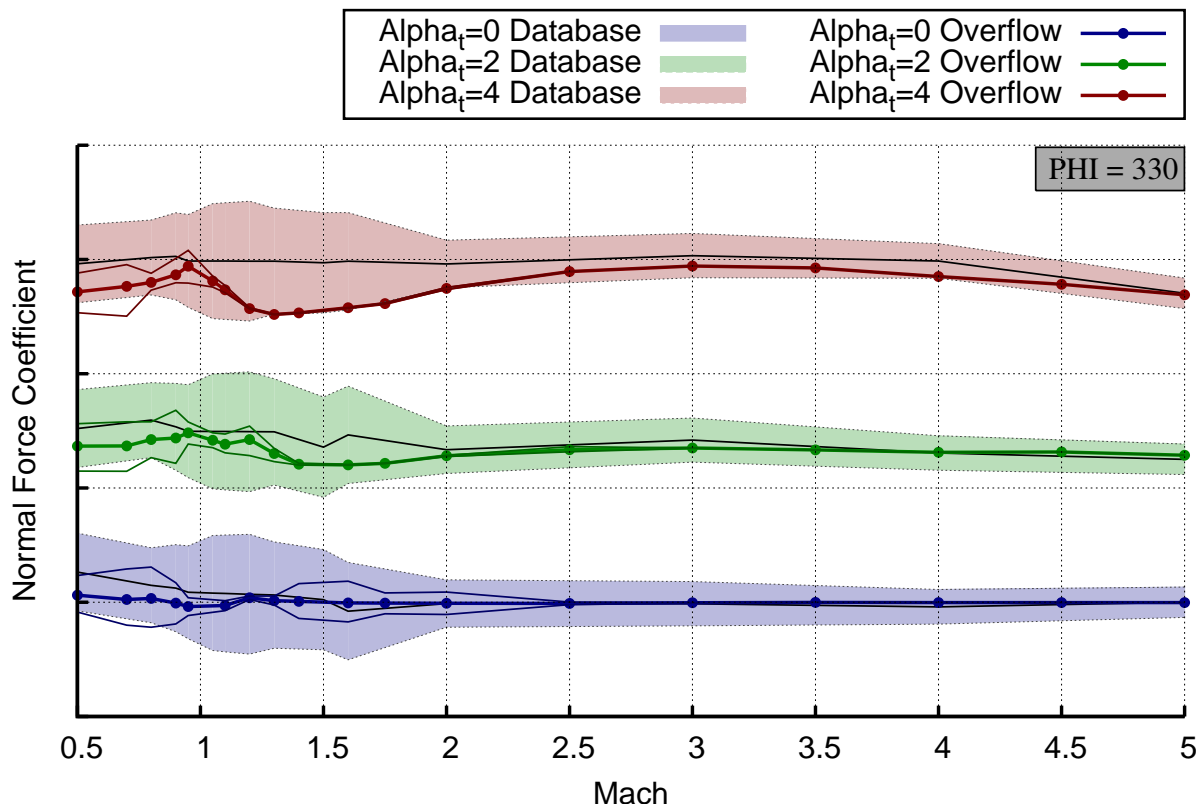


Figure 16. SLS aerodynamic database comparison of CN versus Mach number for  $\phi = 330$

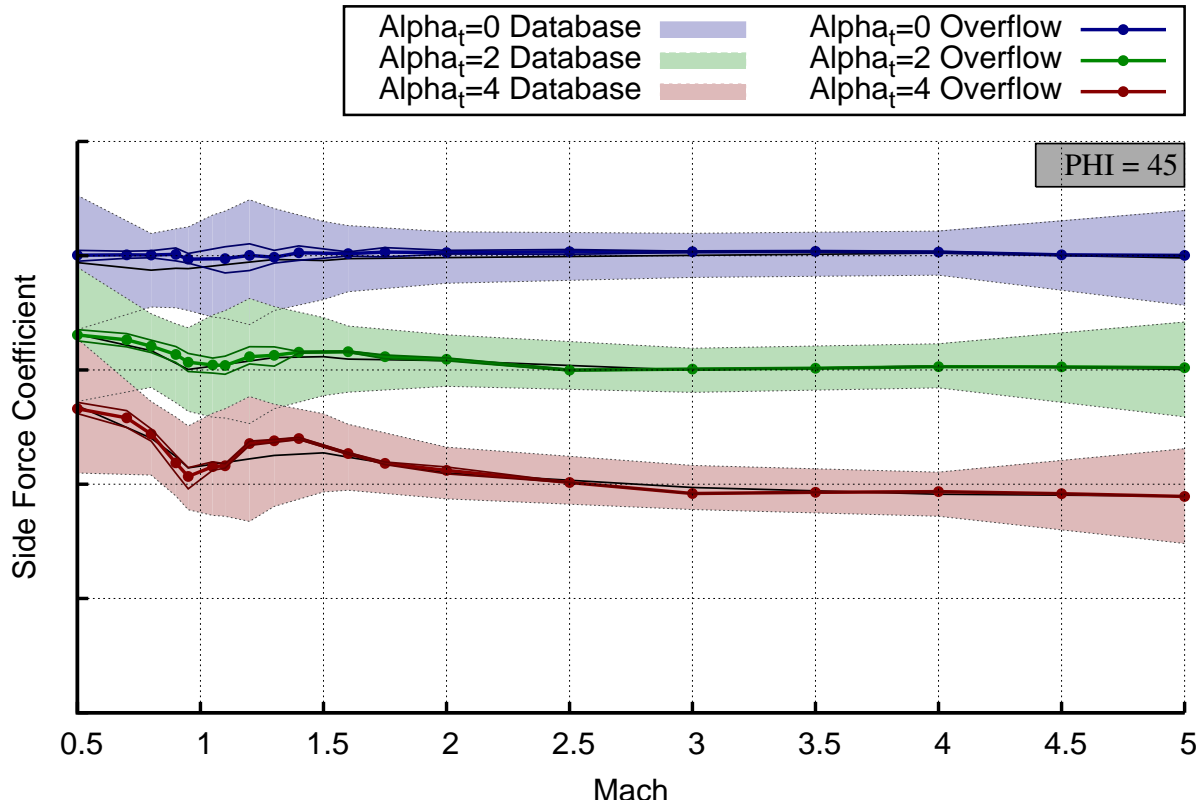


Figure 17. SLS aerodynamic database comparison of CY versus Mach number for  $\phi = 45$

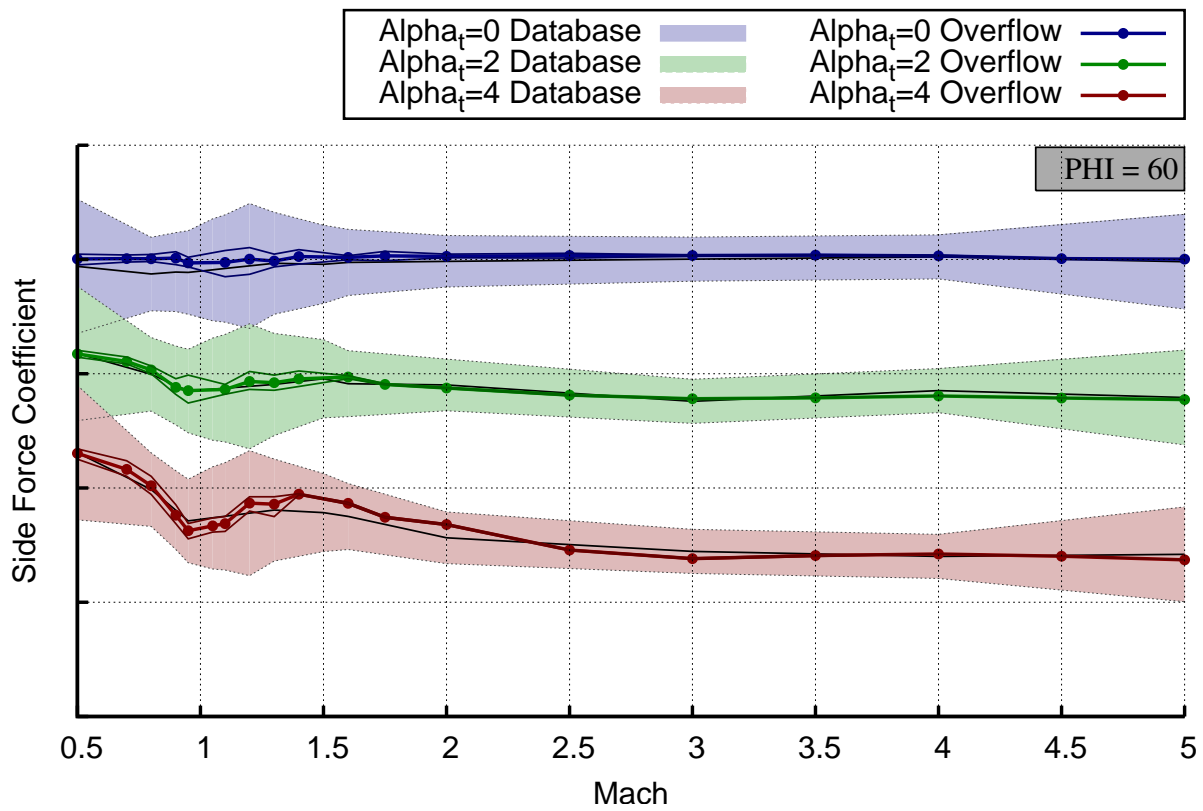


Figure 18. SLS aerodynamic database comparison of CY versus Mach number for  $\phi = 60$

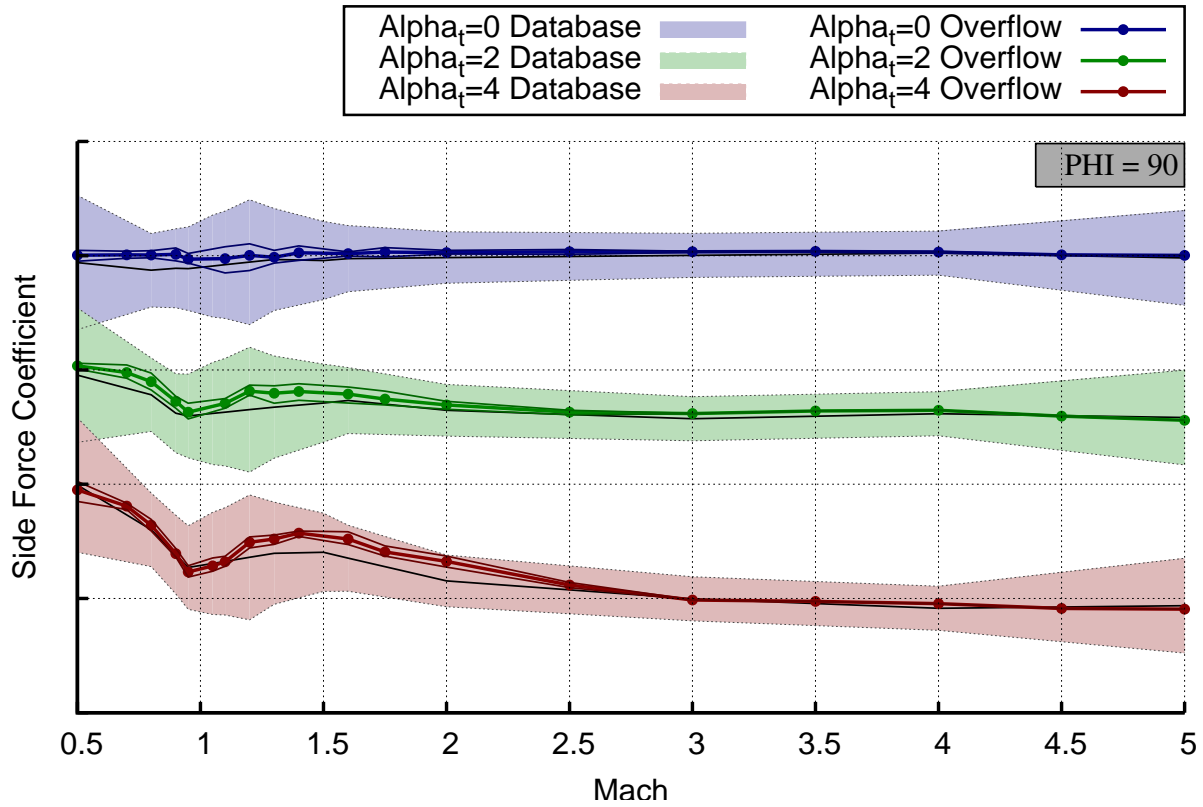


Figure 19. SLS aerodynamic database comparison of CY versus Mach number for  $\phi = 90$

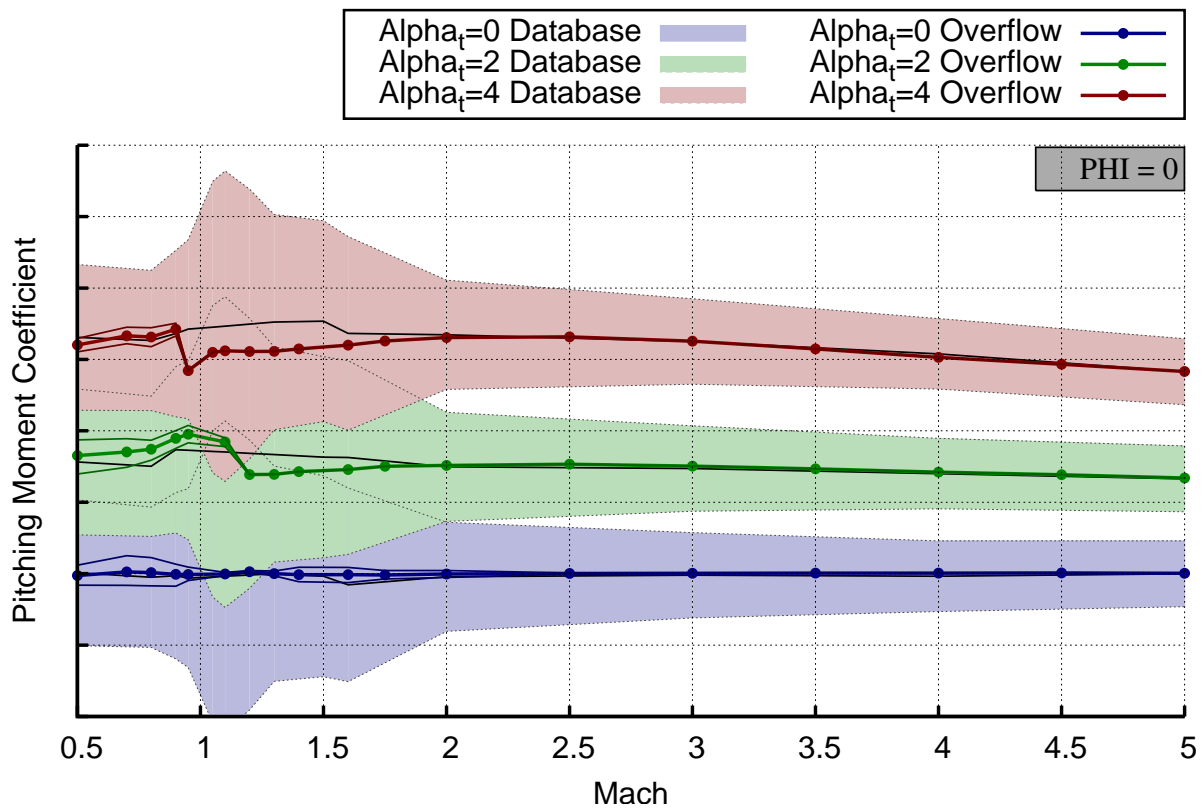


Figure 20. SLS aerodynamic database comparison of CLM versus Mach number for  $\phi = 0$

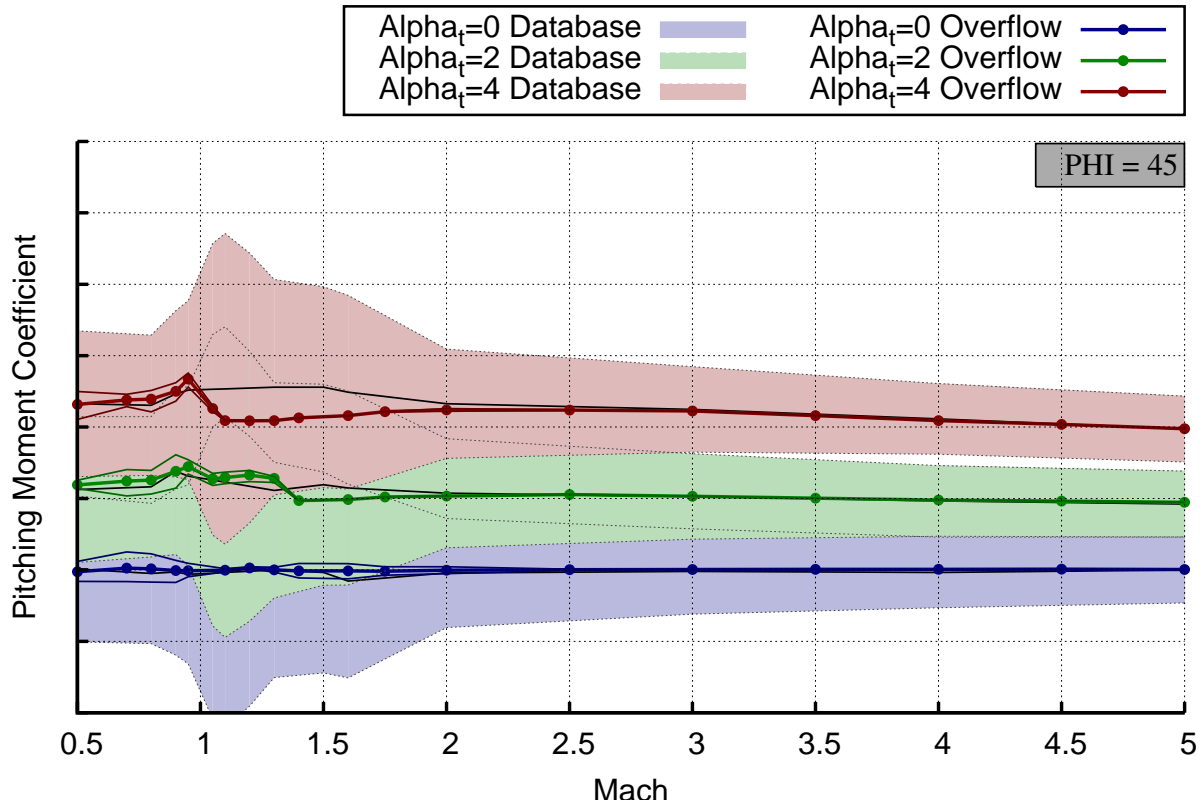


Figure 21. SLS aerodynamic database comparison of CLM versus Mach number for  $\phi = 45$

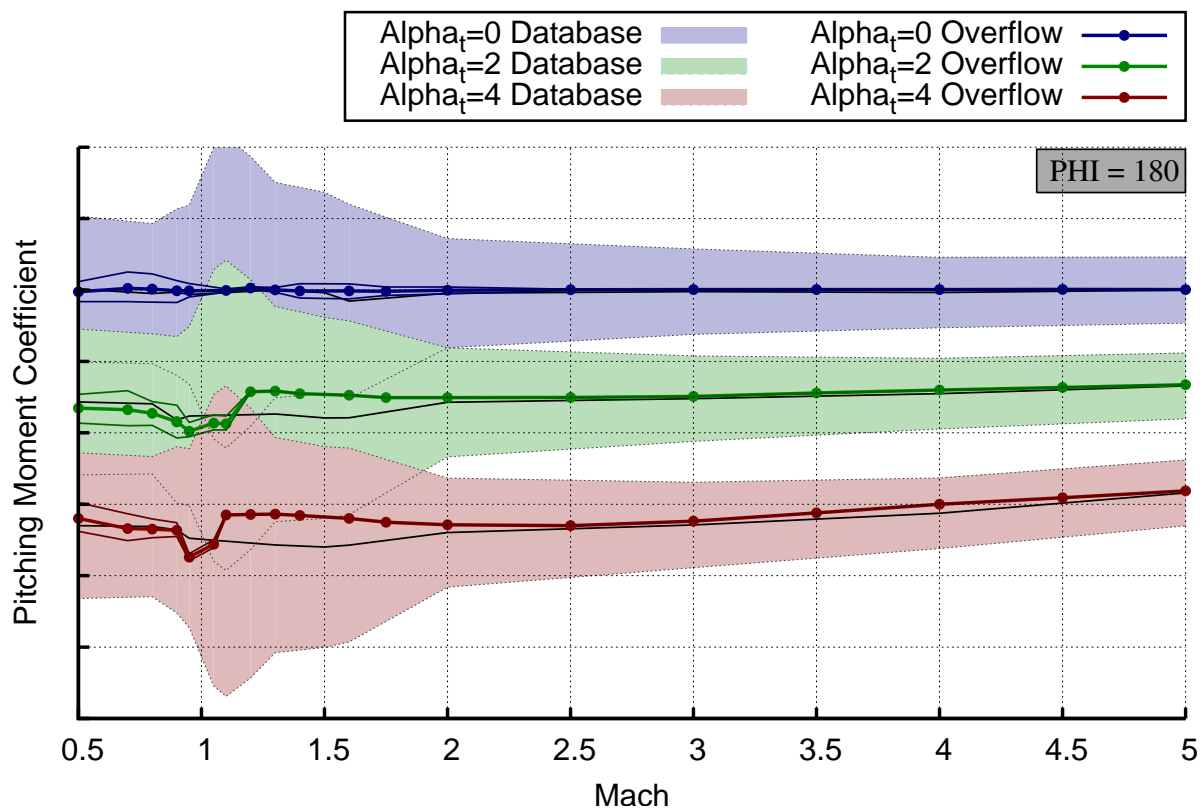


Figure 22. SLS aerodynamic database comparison of CLM versus Mach number for  $\phi = 180$

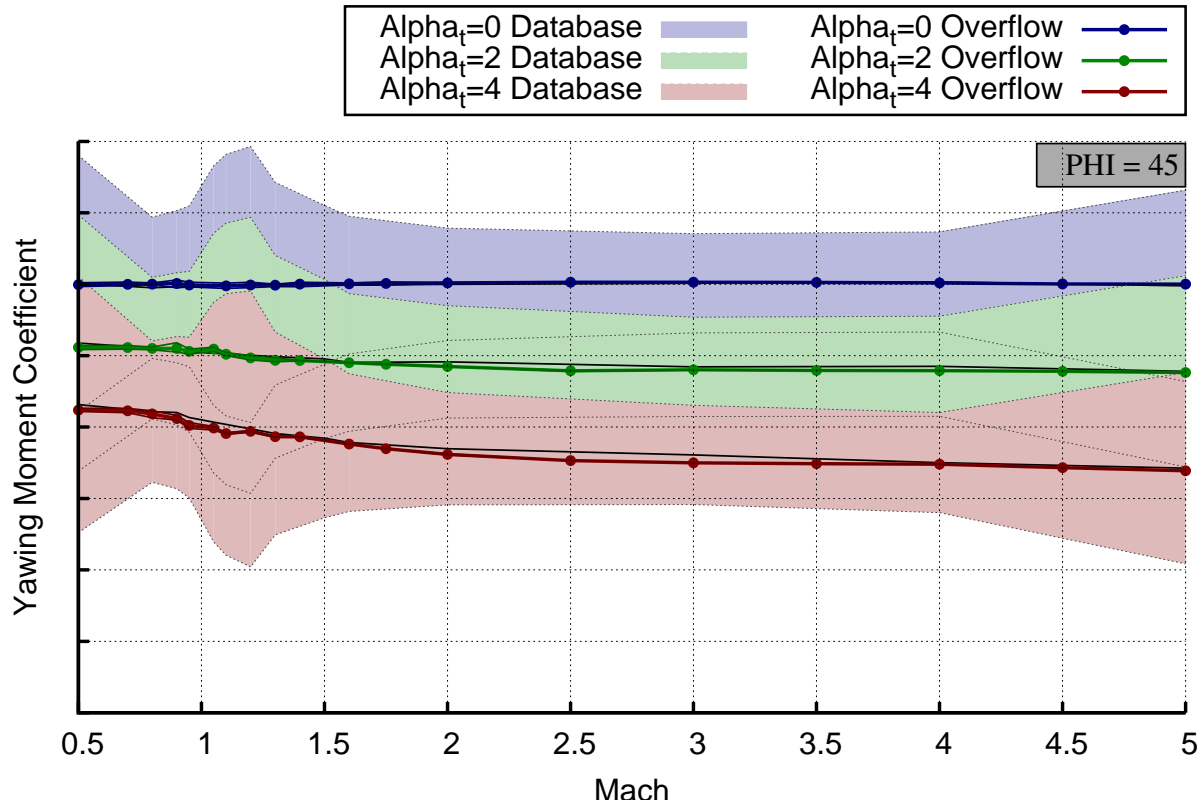


Figure 23. SLS aerodynamic database comparison of CLN versus Mach number for  $\phi = 45$

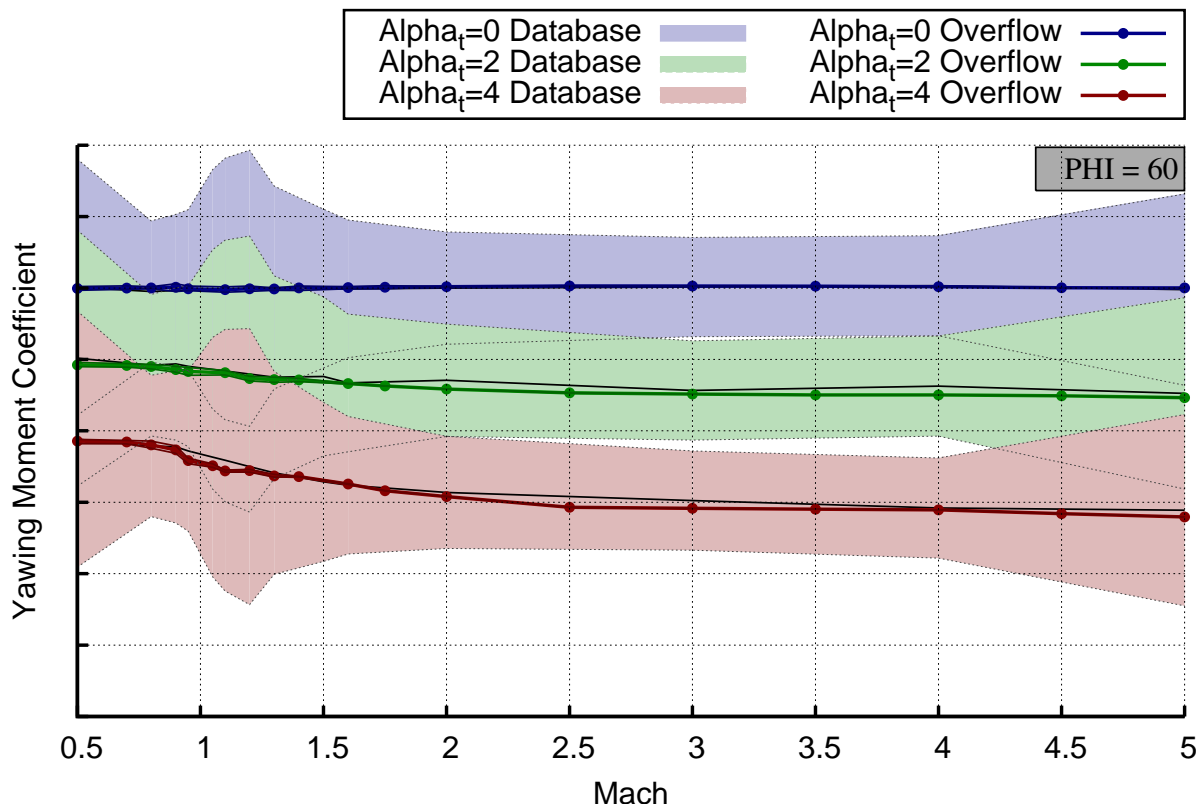


Figure 24. SLS aerodynamic database comparison of CLN versus Mach number for  $\phi = 60$

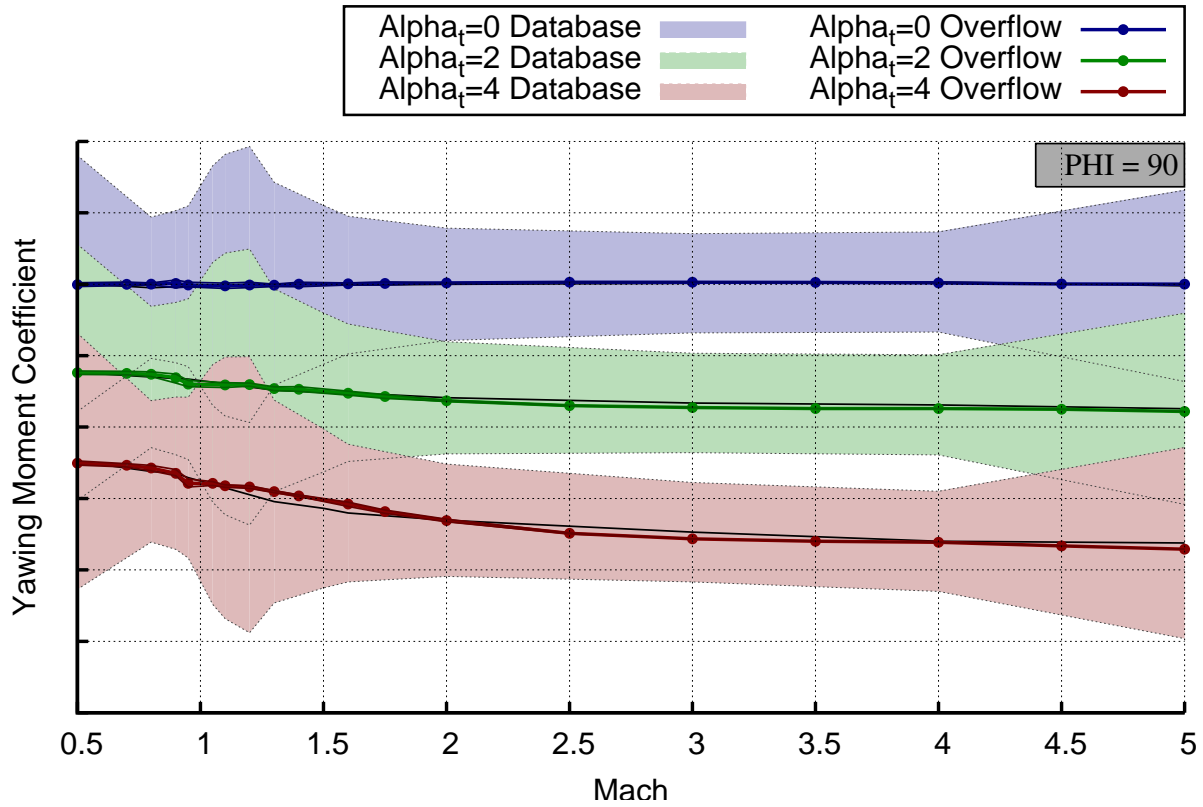


Figure 25. SLS aerodynamic database comparison of CLN versus Mach number for  $\phi = 90$

thus it was not possible to include most of the protuberances that are included in the CFD geometry. While it is outside the scope of the current work to quantify the effect of each of these differences, we can attempt to understand the CFD accuracy by running some cases with different inputs. This topic is explored in the next section.

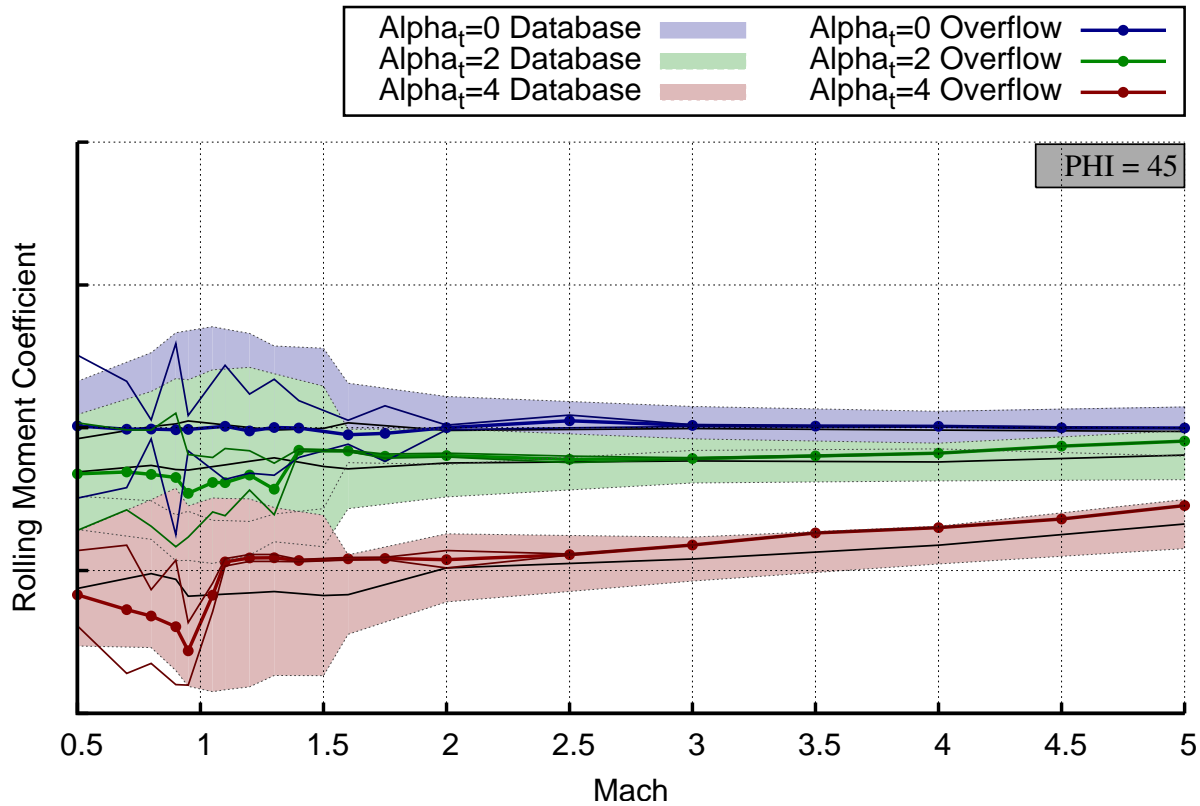


Figure 26. SLS aerodynamic database comparison of CLL versus Mach number for  $\phi = 45$

## X. Sensitivity to Turbulence Model and Differencing Scheme

Four cases were run with two additional sets of inputs to the Overflow code. The first of these was the choice of the SST turbulence model to replace the Spalart-Allmaras turbulence model. The second of these was the choice of the third-order upwind-difference HLLC scheme to replace the central-differencing scheme, together with an SSOR line-relaxation implicit solver to replace the block-tridiagonal approximate factorization implicit solver. These four sensitivity cases were run at flight conditions of  $\alpha_t = 4.0$  degrees,  $\phi = 210.0$  degrees, and Mach numbers of 0.95, 1.1, 1.3, and 1.6. These conditions were chosen because of the differences seen between the CFD and the database at these Mach numbers and at this angle of attack.

Figures 27 through 31 plot the aerodynamic coefficients CN, CY, CLM, CLN, and CLL versus Mach number for  $\phi = 210$  degrees. The results of the sensitivity runs are included as dark-gray and light-gray circles, representing the SST-model results and the HLLC upwind results, respectively. These figures show the differences between the baseline runs and the sensitivity runs in the vehicle integrated forces and moments. The differences relative to the aerodynamic database uncertainty values are very small, and they are generally smaller than the difference between the time-averaged CFD results and the nominal aerodynamic database values. When comparing the local surface pressure coefficients ( $C_p$ ), however, the differences between the CFD solutions are more significant. Figures 32 through 35 show surface  $C_p$  on the vehicle from the baseline solutions and from the SST solutions. Positioned in between are surface contour plots of  $\Delta C_p$ , calculated as the baseline  $C_p$  minus the SST-solution  $C_p$ . The same comparisons are made for upwind-differencing solutions. It can be seen that the biggest differences in the CFD solutions occur in regions of

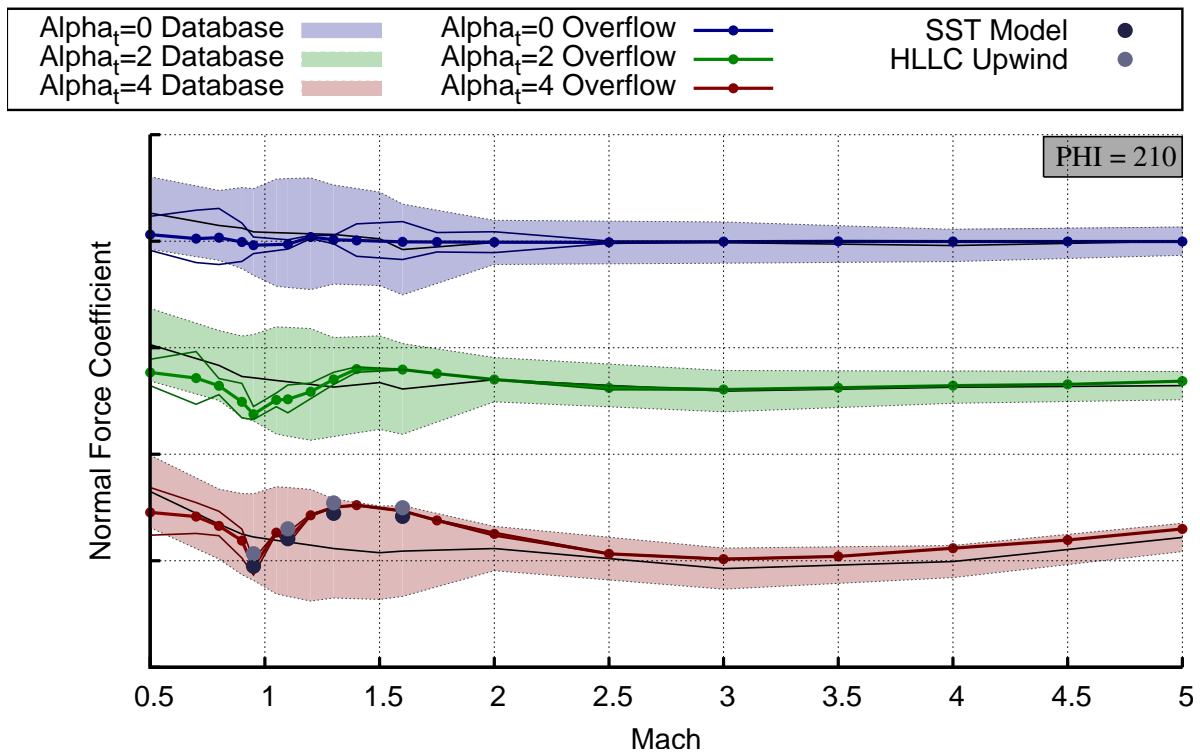


Figure 27. SLS aerodynamic database comparison of CN versus Mach number for  $\phi = 210$

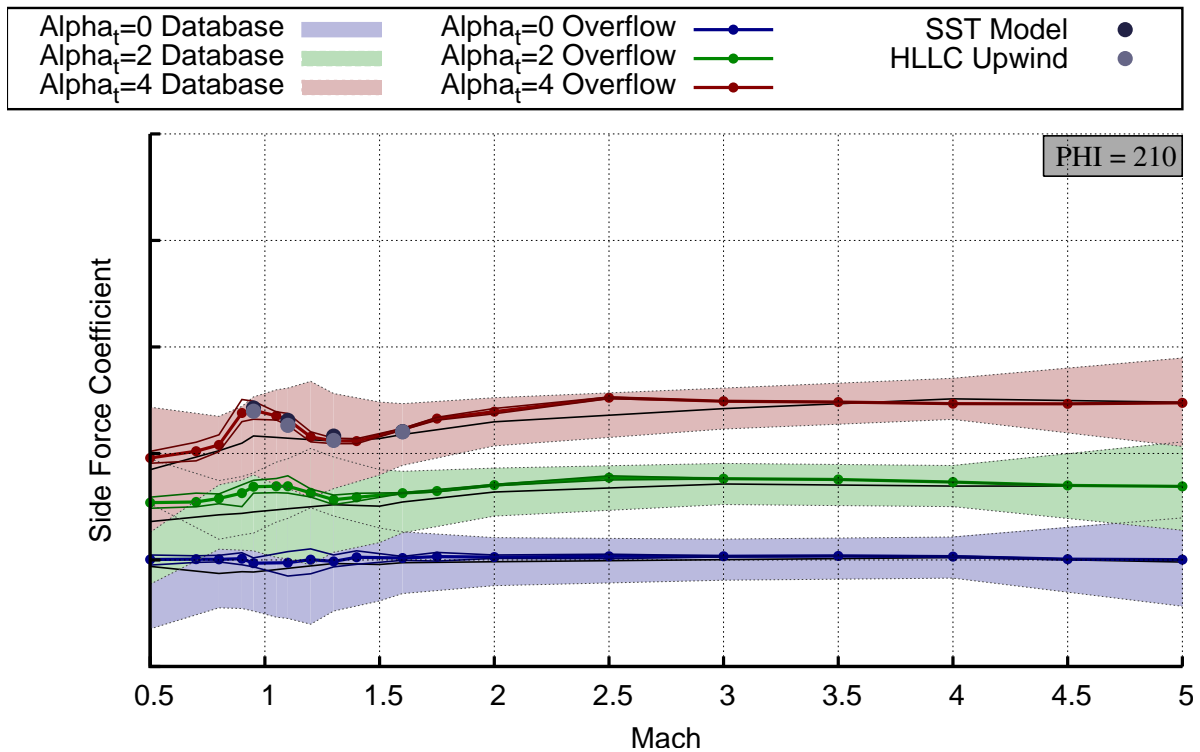


Figure 28. SLS aerodynamic database comparison of CY versus Mach number for  $\phi = 210$

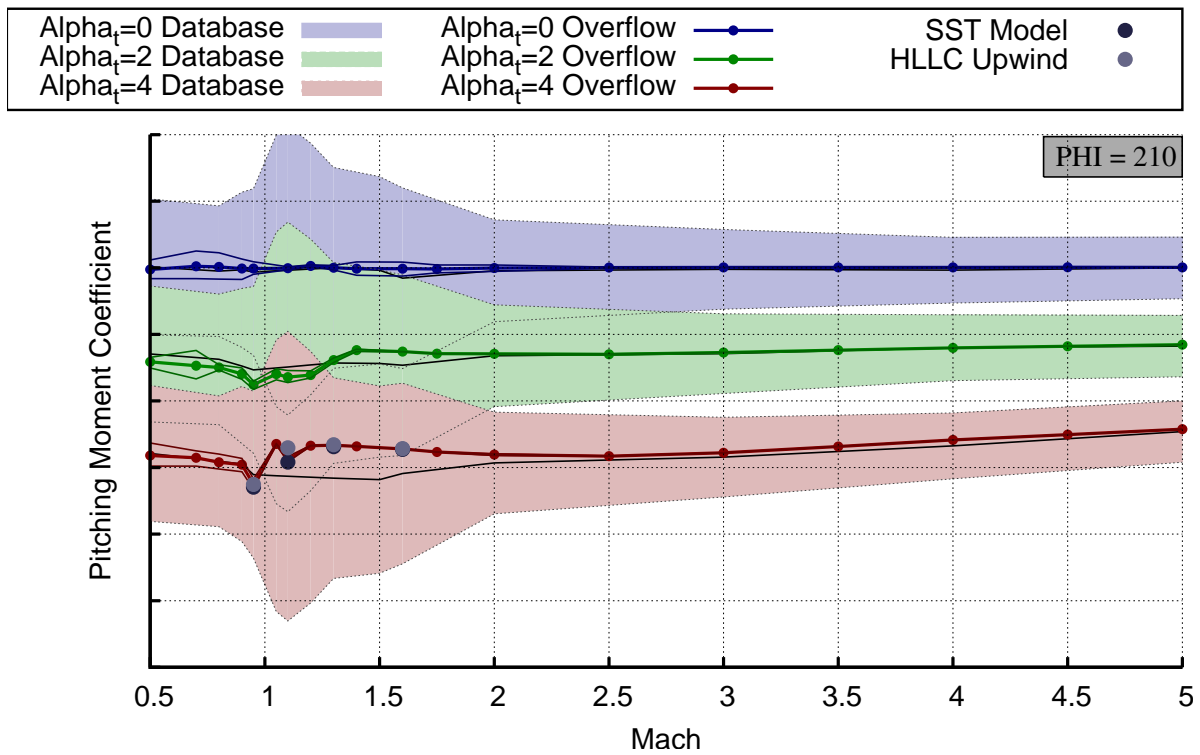


Figure 29. SLS aerodynamic database comparison of CLM versus Mach number for  $\phi = 210$

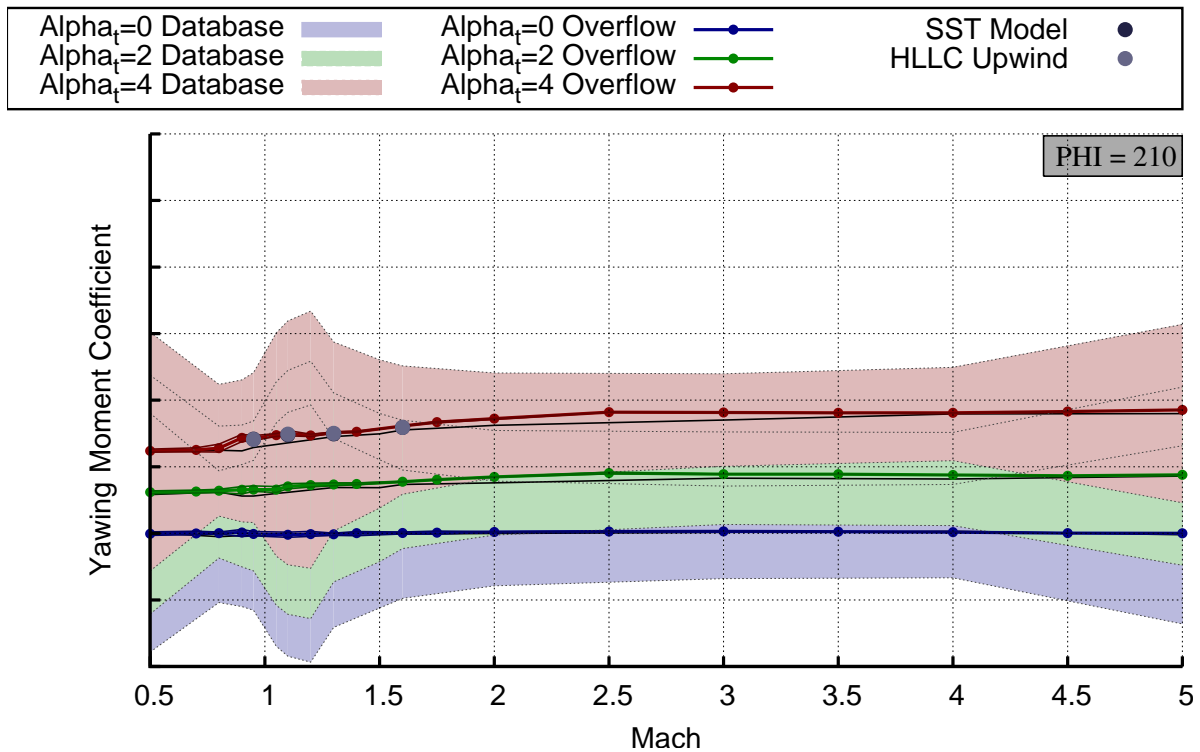


Figure 30. SLS aerodynamic database comparison of CLN versus Mach number for  $\phi = 210$

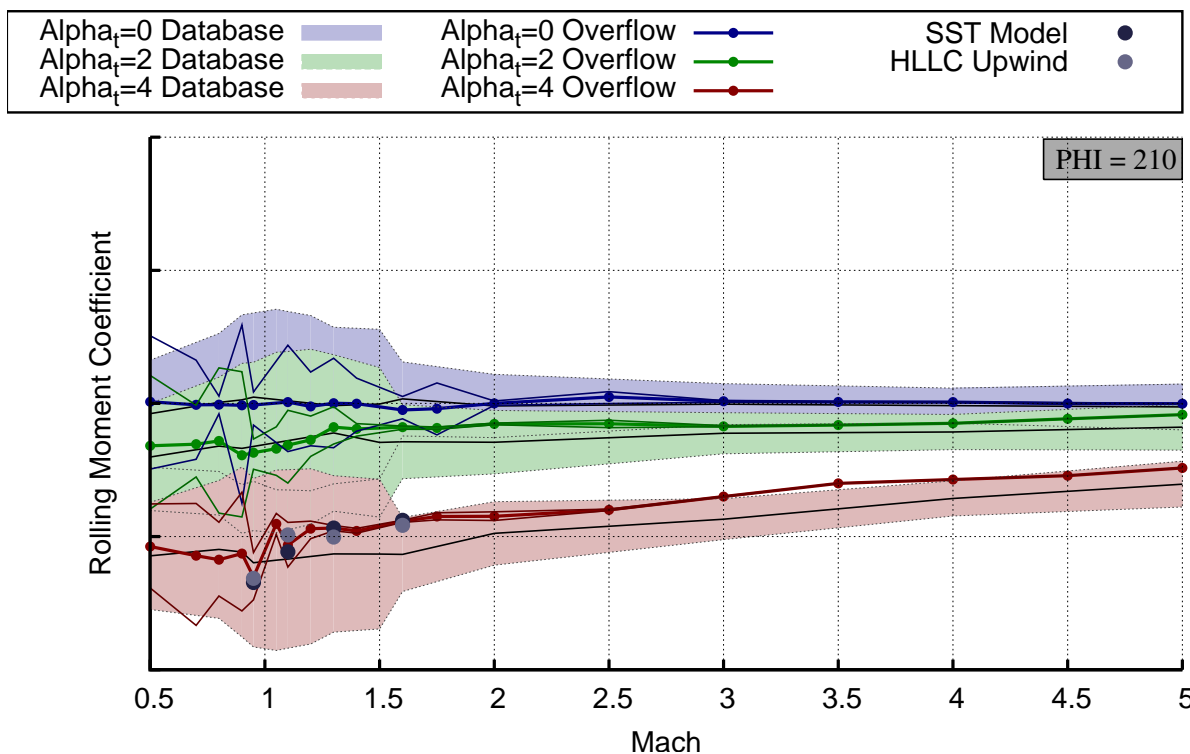


Figure 31. SLS aerodynamic database comparison of CLL versus Mach number for  $\phi = 210$

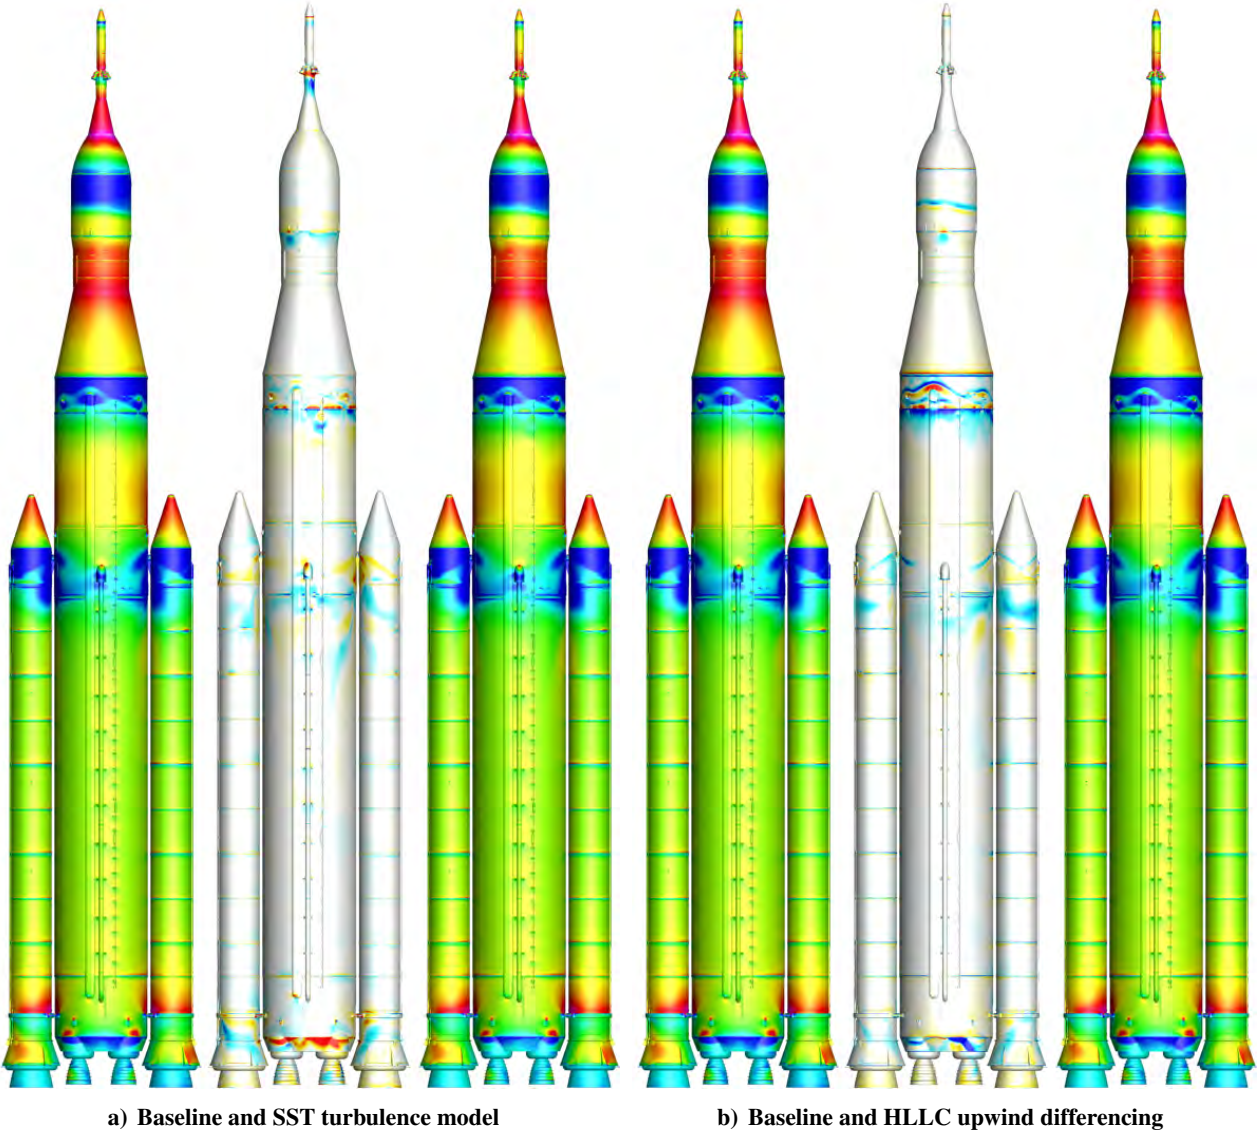
shock interactions and regions of flow separation from the surface of the vehicle.

Comparisons of the sectional loads in the normal and lateral directions on the core and SRBs are shown in Fig. 36 for  $M = 1.1$ , and in Fig. 37 for  $M = 1.6$ . The  $M = 1.1$  differences in the loads between the cases are substantial, especially for the lateral loads on the core. The biggest differences seen in the SST turbulence-model solutions appear to be in regions with flow separation, for example, near the forward and aft attach hardware. The biggest differences seen in the HLLC upwind-scheme solutions appear to be in high-pressure gradient regions caused by shock waves.

## XI. Booster Separation

The SLS launch vehicle includes two solid-rocket boosters that burn out before the primary core stage and thus must be discarded during the ascent trajectory. Because gravity is not acting in a way that pulls the expended boosters away from the core, and aerodynamic loads also do not act in a way to substantially encourage separation, small rockets called Booster Separation Motors (BSMs) are used to actively push the boosters away from the core. There are 16 of these BSMs—four forward BSMs and four aft BSMs on each booster. The forward BSMs fire in a partially forward (upstream-facing) direction, which results in complex plumes and interactions of those plumes with the high-speed freestream flow, and with the core stage of the vehicle. In addition, the low dynamic pressure during the separation event combined with the relatively high-thrust separation motors means that the plume affects a very large volume of flow on the scale of the SLS vehicle itself. To top this all of this off, the flow is unsteady. In addition to the complex flow physics, the biggest challenge for creating a booster-separation aerodynamic database are the large number of basis variables, including orientation of the core, relative position and orientation of the boosters, and rocket thrust levels.

The design of the SRBs is relatively unchanged from that used with the Space Shuttle Launch Vehicle, which has a rich flight history. However, the difference between the SLS core and the Space Shuttle External Tank results in much narrower clearances during SRB separation. In addition, the placement of the main



**Figure 32. Surface  $C_p$  and  $\Delta C_p$  at Mach 0.95**

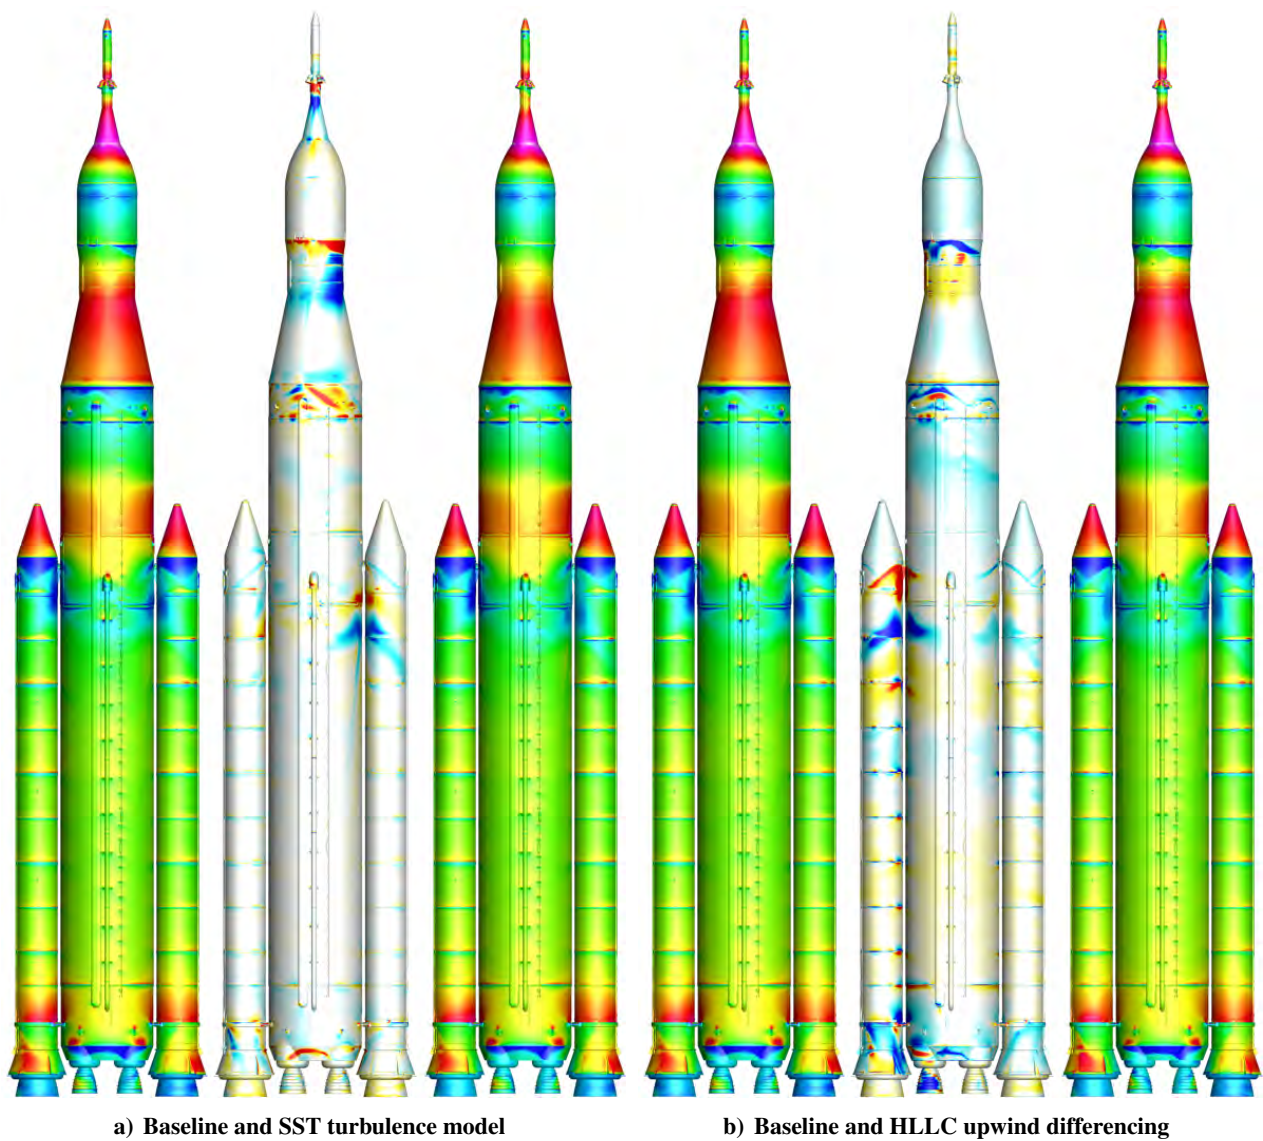


Figure 33. Surface  $C_p$  and  $\Delta C_p$  at Mach 1.1

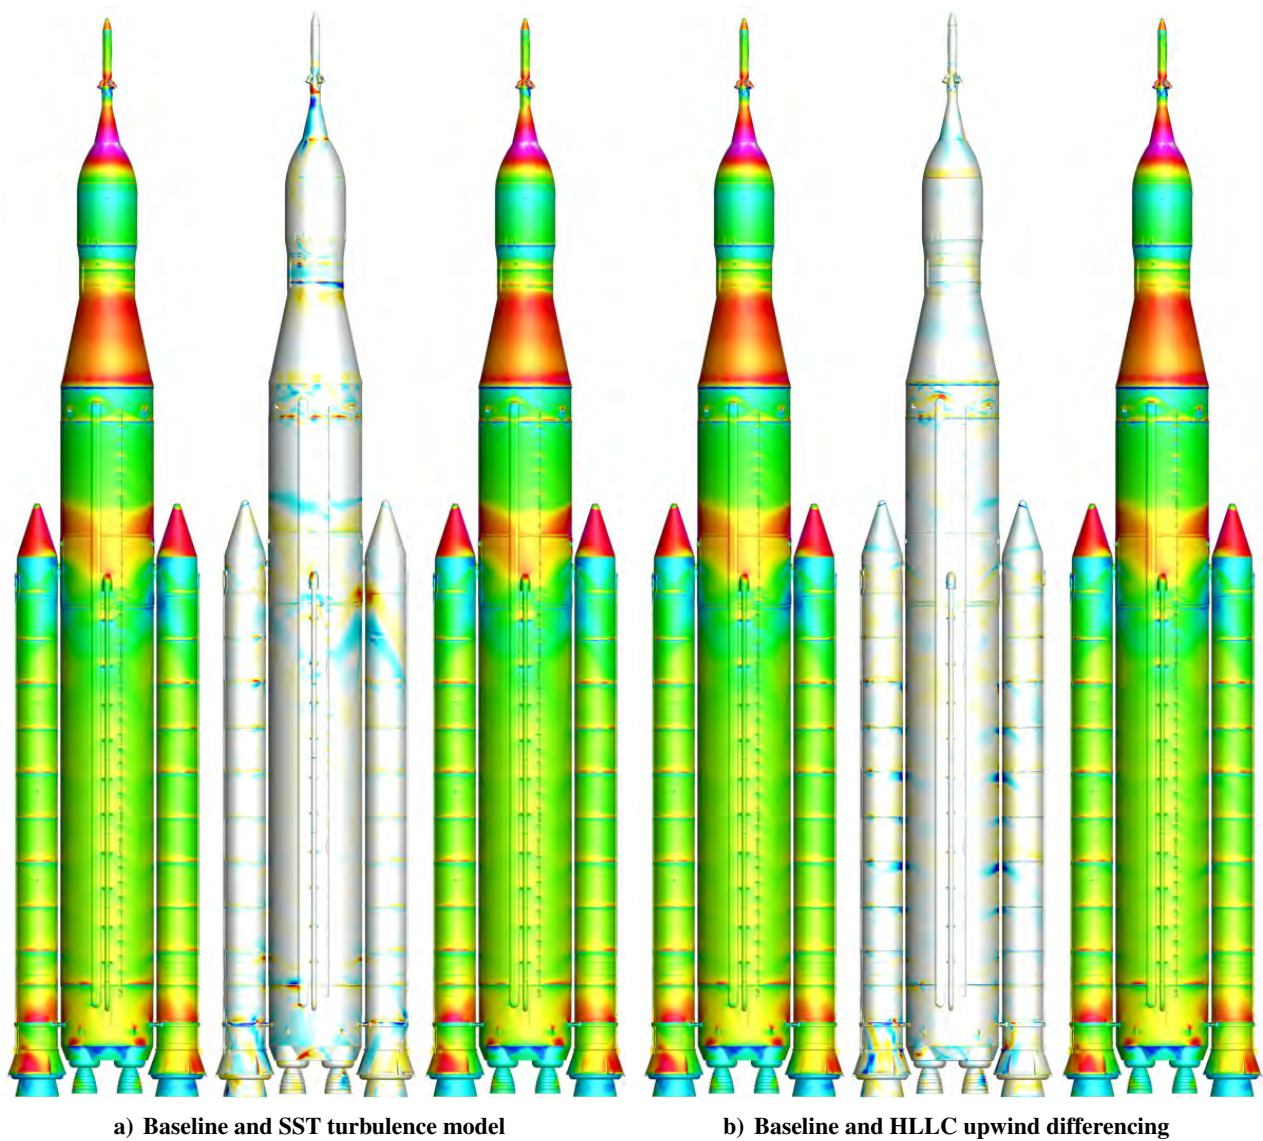


Figure 34. Surface  $C_p$  and  $\Delta C_p$  at Mach 1.3

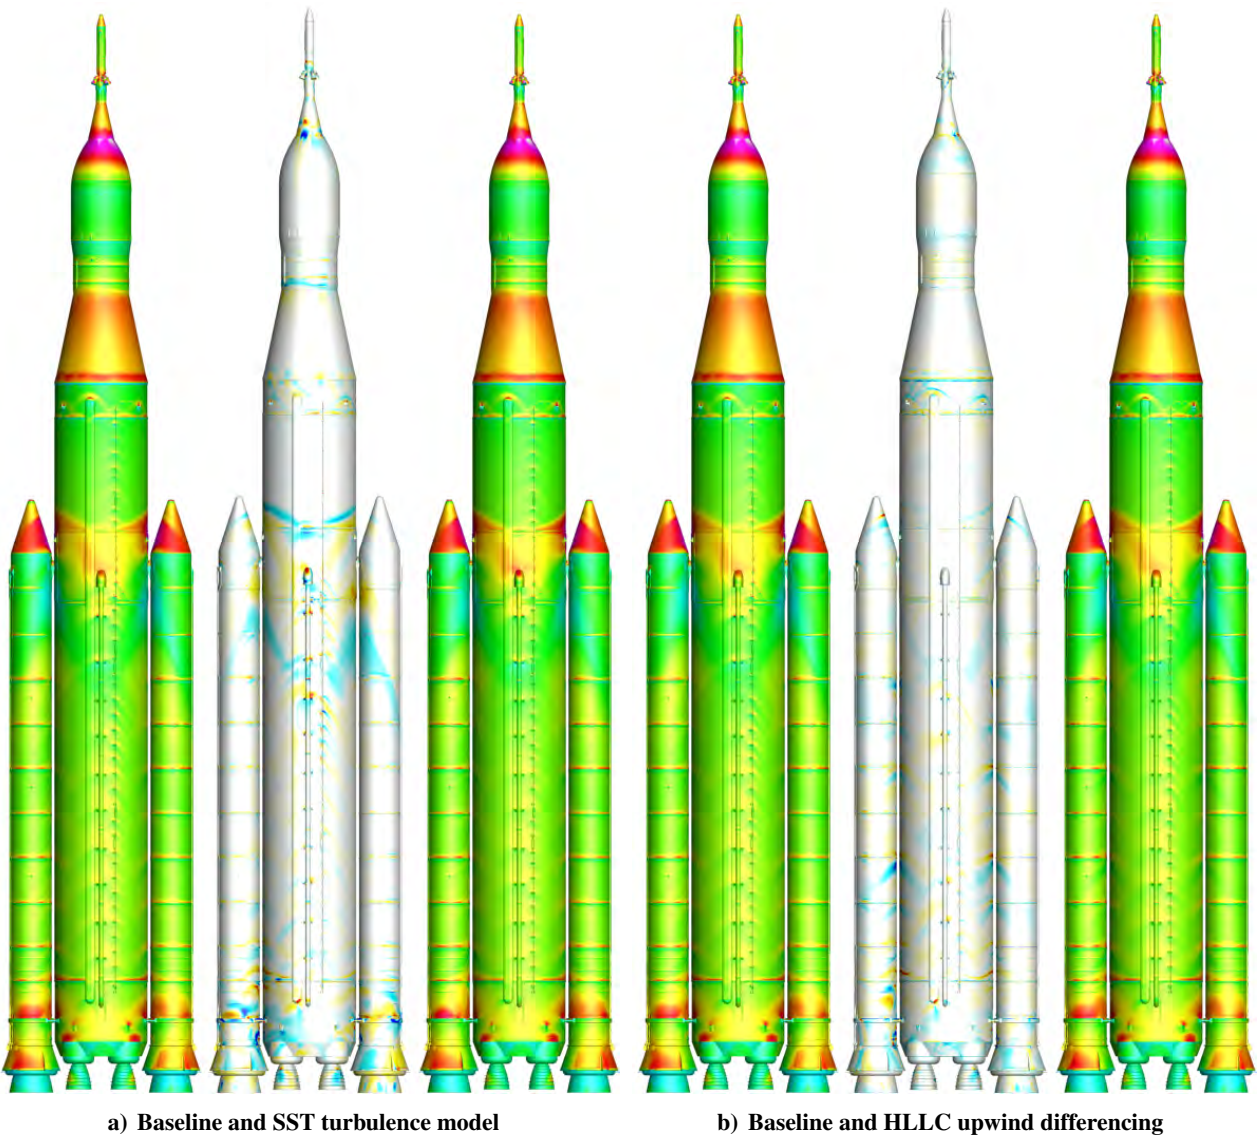
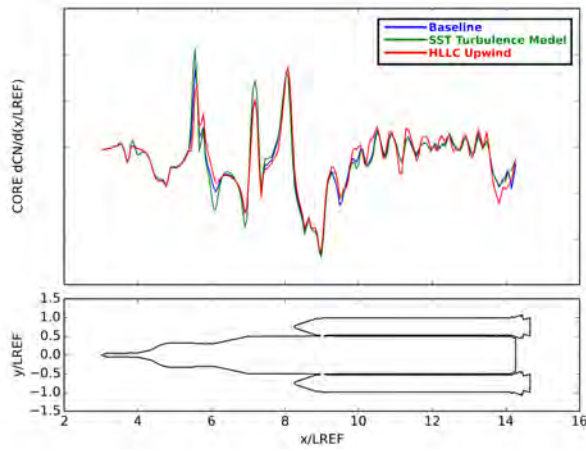
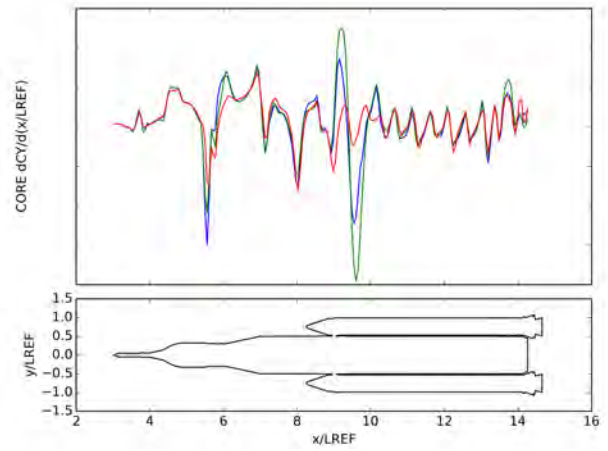


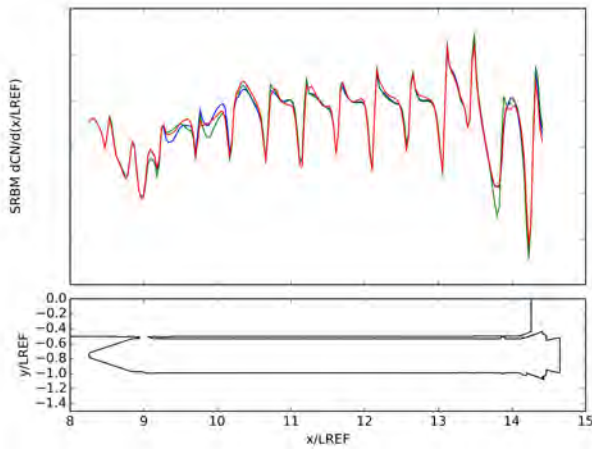
Figure 35. Surface  $C_p$  and  $\Delta C_p$  at Mach 1.6



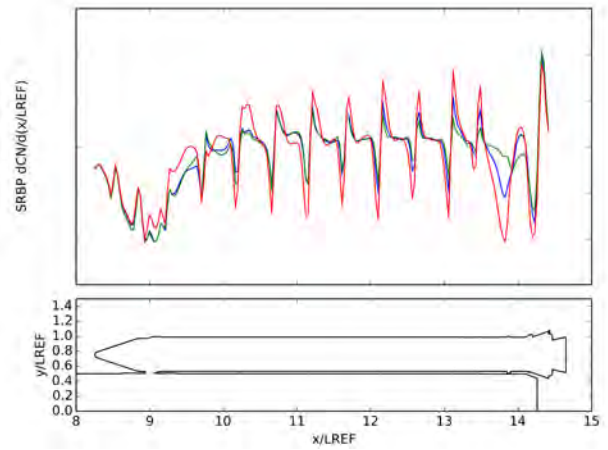
a) Normal sectional loads on core



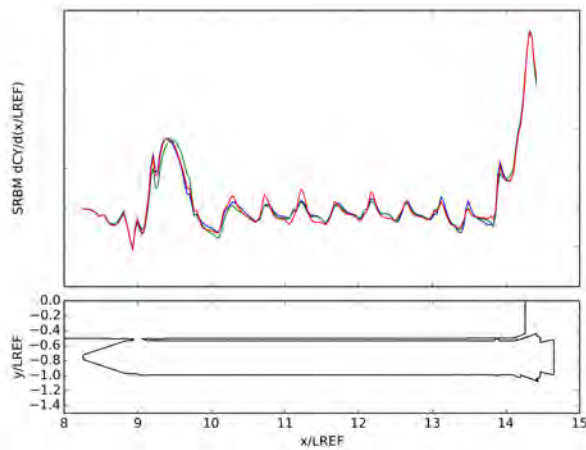
b) Lateral sectional loads on core



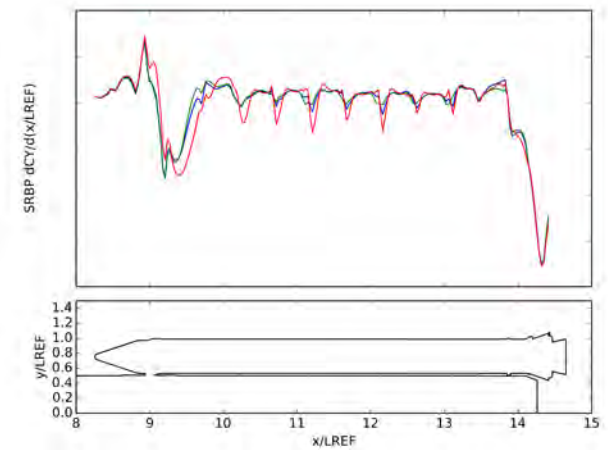
c) Normal sectional loads on left booster (SRBM)



d) Normal sectional loads on right booster (SRBP)

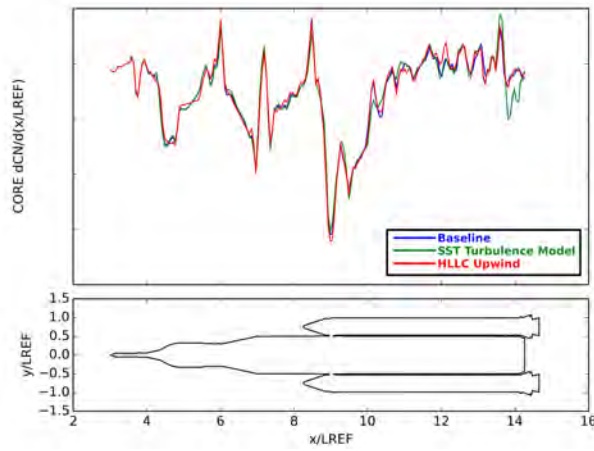


e) Lateral sectional loads on left booster (SRBM)

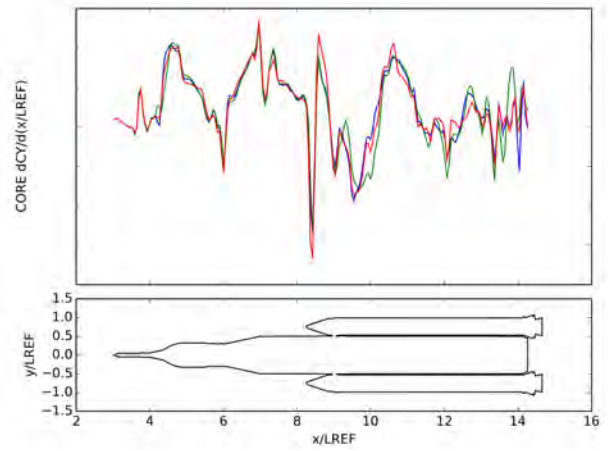


f) Lateral sectional loads on right booster (SRBP)

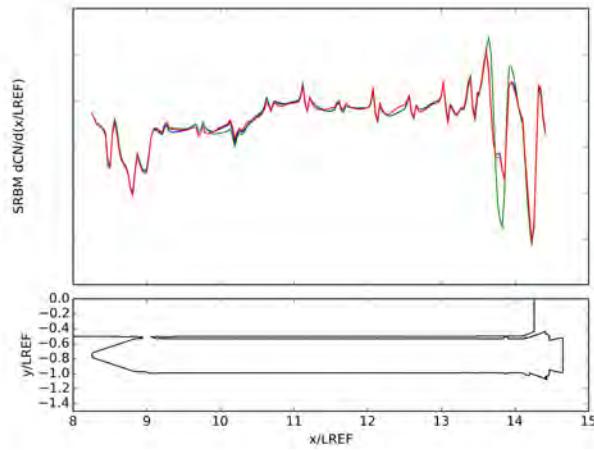
**Figure 36. Sectional aerodynamic loads for turbulence-model and differencing-scheme sensitivity cases on the solid rocket boosters at Mach 1.1**



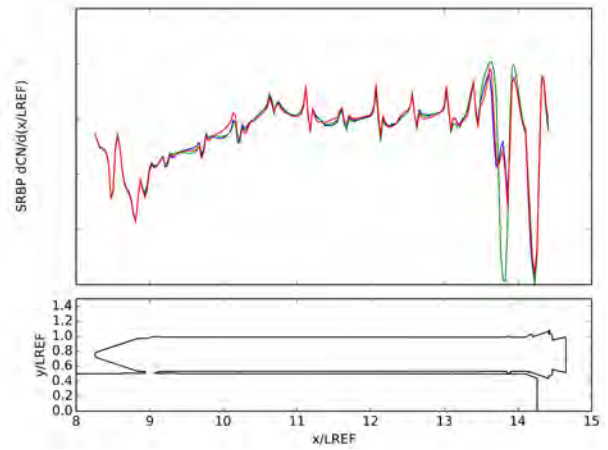
a) Normal sectional loads on core



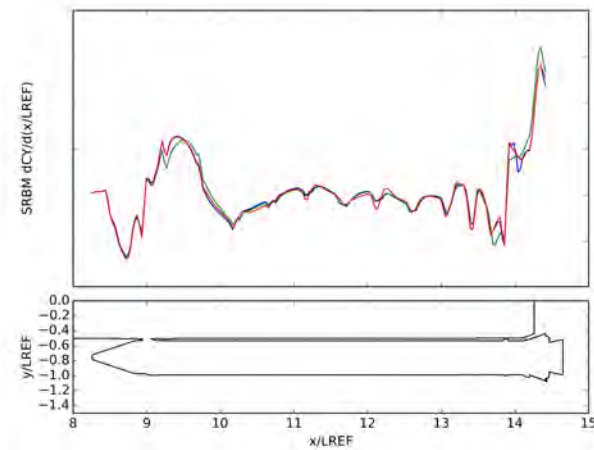
b) Lateral sectional loads on core



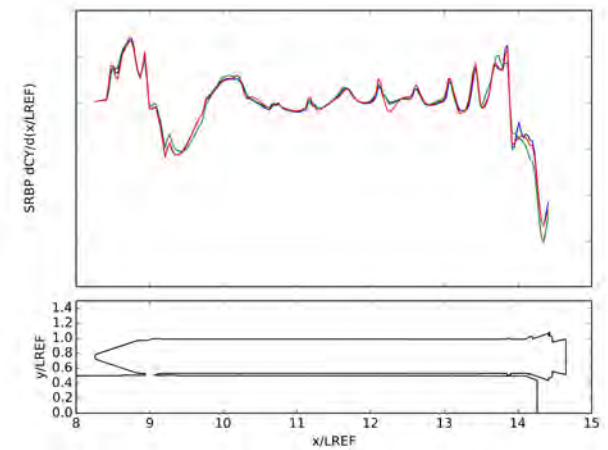
c) Normal sectional loads on left booster (SRBM)



d) Normal sectional loads on right booster (SRBP)



e) Lateral sectional loads on left booster (SRBM)



f) Lateral sectional loads on right booster (SRBP)

**Figure 37. Sectional aerodynamic loads for turbulence-model and differencing-scheme sensitivity cases on the solid rocket boosters at Mach 1.6**

engines at the base of the core results in the main-engine plumes being directly in the path of the separating SRBs as they move away from the core, causing a plume interaction that was not present during separation from the Space Shuttle Launch Vehicle. Because of these additional challenges, it is necessary to reduce aerodynamic uncertainty as much as possible in order to clear the integrated system for flight.

Because of the very large run matrix required to create a database for the range of possible configurations that occur during separation, an inviscid flow solver with output-based adaptive meshing called Cart3D [19, 20] was selected as the primary analysis tool. This tool requires only a surface grid and flow solver inputs to run; volume mesh creation is automatic. Combined with the adaptive mesh capability, this means that quality solutions can be obtained with only minimal oversight by the user. Figure 38 shows an example of the mesh adaptation results for a case soon after separation, in which the BSMs are firing, and the SRBs are close to the core. The initial mesh, in Fig. 38a, has 416 thousand volume cells, and the final mesh has 13 million. A mesh growth ratio of 1.5 was used for the first two adaptation cycles, and the remaining cycles used a growth ratio of 2.0. This calculation took 1.5 hrs using 20 CPUs on a single Ivy Bridge node.

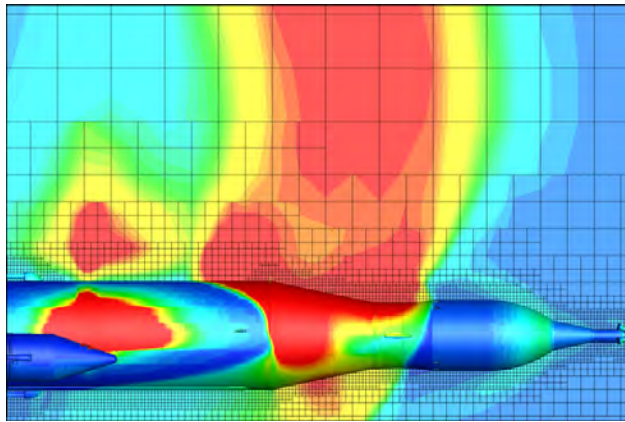
The forward BSMs cause a strong shock to form as far upstream as the MPCV service module, and the flow is subsonic behind this shock. The mesh adaptation history demonstrates how relatively high resolution is required quite far off of the body. Thus a single grid able to resolve all flow conditions in the database would have an excessively high number of cells.

To support the SLS project, a new booster-separation aerodynamic database is currently being produced. Comparisons will be made to viscous Overflow results, which will provide insights that are hard to predict for a database with complex flow and many degrees of freedom. The procedures for this work will follow the lessons learned to perform first-stage separation analysis that supported Ares I [4, 21]. In addition, a wind tunnel study performed at NASA Langley Research Center's Unitary Plan Wind Tunnel in September and October of 2014 will be used to provide validation for the computational aerodynamic database. Because of the range of geometric scales, from booster separation motors with throat planes several inches in diameter to the overall vehicle length of more than 300 ft, completing a wind tunnel test to support the aerodynamic database was judged to be prohibitively expensive. As a result, the CFD database generated in this work is a high-priority analysis for the mission success of SLS.

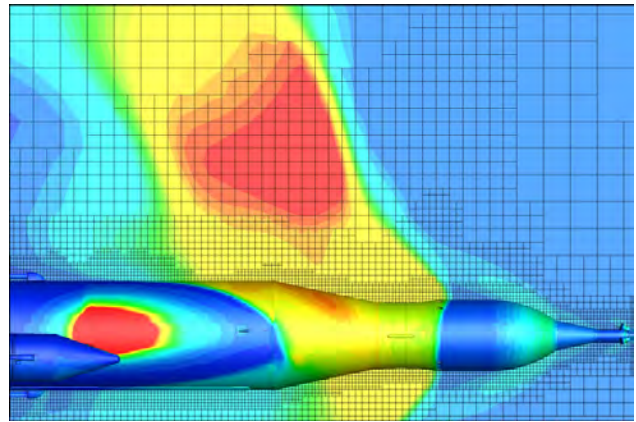
## Conclusions

The current CFD effort has delivered valuable data products to the SLS program in support of the development of the new space-launch capability. Over 28 million CPU hours were used to run Overflow simulations for 739 flight conditions covering a wide range of ascent Mach numbers, total angles of attack, and roll angles. CFD derived data has been delivered to the program in the form of sectional line loads, surface  $C_p$ , vent-location pressures and boundary-layer data, and aerodynamic loads on all of the vehicle's protuberances. Favorable comparisons were found between the CFD flight-vehicle aerodynamics and the wind-tunnel derived aerodynamic database. The regions of largest discrepancy occur in the transonic Mach range, a flight regime with complex shock interactions and unsteady flows, and a range which presents significant challenges to the wind-tunnel testing. Four of the cases were run to evaluate the solutions using a different turbulence model, and also with an upwind differencing scheme. The differences in the integrated forces and moments are relatively small, but there are larger local differences in  $C_p$  in regions with shock waves and separated flow.

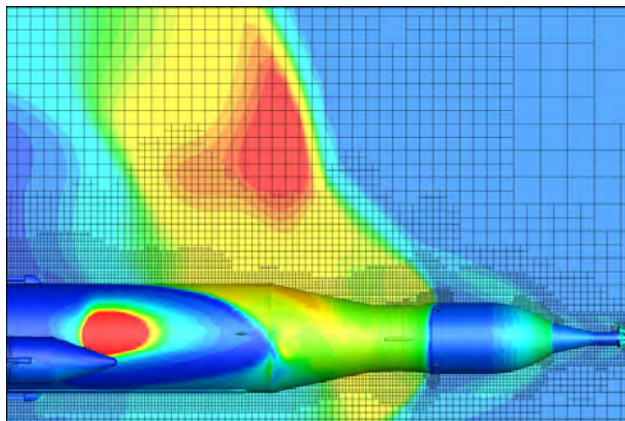
Analysis of the aerodynamics during booster separation using the Cart3D unstructured Cartesian solver is also being performed for the program. These simulations benefit greatly from the use of output-based adaptive meshing. The complex flow physics generated by 16 booster-separation motors firing into the supersonic flow-field requires resolution of many off-body flow features. This analysis work will continue in the near future as the program builds an aerodynamic database which will be used to ensure successful booster separation. The large number of basis variables will require tens of thousands of simulations. The accuracy of the simulations will be assessed by comparison to newly acquired wind-tunnel data, with the goal of reducing the aerodynamic uncertainty within the database.



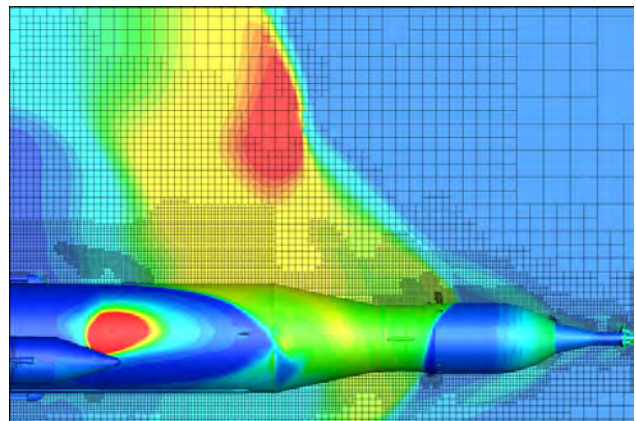
a) Initial mesh, 0 adaptation cycles



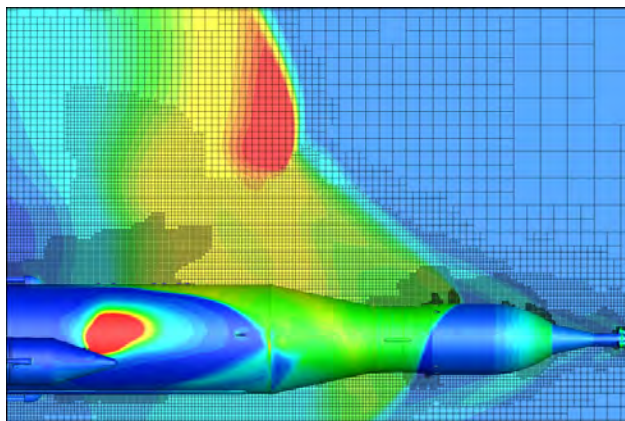
b) After 2 adaptation cycles



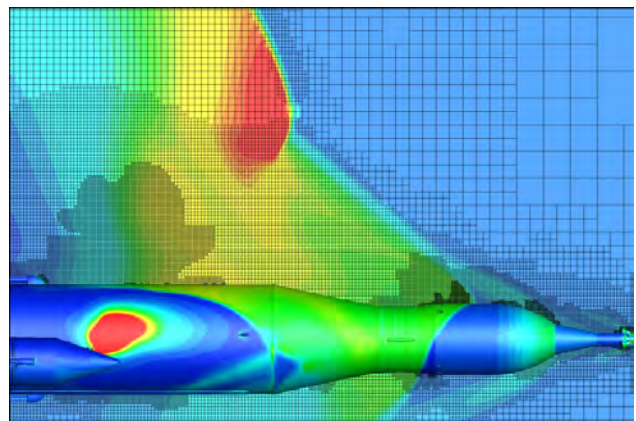
c) After 3 adaptation cycles



d) After 4 adaptation cycles



e) After 5 adaptation cycles



f) After 6 adaptation cycles

Figure 38. Grid history of static pressure on vehicle surface and center  $xz$ -plane cut for case with BSMs firing

## Acknowledgments

This current work was made possible by the the hard work of many individuals. The authors wish to acknowledge the following individuals for their contributions to defining the initial geometry, to the grid generation, and to the flow-solver inputs. Thanks go to Jeffrey Housman, William Chan, Emre Sozer, Shayan Moini-Yekta, H. Dogus Akaydin, Jeff Onufer, Goetz Klopfer, and Cetin Kiris. We wish to acknowledge the expertise and hard work of the people of the NASA Advanced Supercomputing Division. Their world-class facility made it relatively easy to utilize millions and millions of CPU hours.

## References

- [1] Gomez, R. J., Vicker, D., Rogers, S. E., Aftosmis, M. J., Chan, W. M., Meakin, R. L., and Murman, S., “STS-107 Investigation Ascent CFD Support,” *34th AIAA Fluid Dynamics Conference*, June 2004, AIAA Paper 2004-2226.
- [2] Murman, S. M., Aftosmis, M. J., and Rogers, S. E., “Characterization of Space Shuttle Ascent Debris Using CFD Methods,” *43rd AIAA Aerospace Sciences Meeting and Exhibit*, June 2005, AIAA Paper 2005-1223.
- [3] Tejnil, E. and Rogers, S. E., “CFD Assessment of Forward Booster Separation Motor Ignition Overpressure on ET XT 718 Ice/Frost Ramp,” *50th AIAA Aerospace Sciences Meeting*, January 2012, AIAA Paper 2012-0673.
- [4] Pamadi, B. N., Pei, J., Pinier, J. T., Holland, S. D., and Covell, P. F., “Aerodynamic Analyses and Database Development for Ares I Vehicle First-Stage Separation,” *Journal of Aircraft*, Vol. 49, No. 5, 2012, pp. 864–874.
- [5] Kiris, C., Housman, J., Gusman, M., Schauerhamer, D., Deere, K., Elmilgui, A., Abdol-Hamid, K., Parlette, E., Andrews, M., and Blevins, J., “Best Practices for Aero-Database CFD Simulations of Ares V Ascent,” *49th AIAA Aerospace Sciences Meeting*, January 2011, AIAA Paper 2011-16.
- [6] Gusman, M., Housman, J., and Kiris, C., “Best Practices for CFD Simulations Launch Vehicle Ascent with Plumes - OVERFLOW Perspective,” *49th AIAA Aerospace Sciences Meeting*, January 2011, AIAA Paper 2011-1054.
- [7] Aftosmis, M. A. and Rogers, S. E., “Effects of Jet-Interaction on Pitch-Control of a Launch Abort Vehicle,” *46th AIAA Aerospace Sciences Meeting and Exhibit*, January 2008, AIAA Paper No. 2008-1281.
- [8] Childs, R. E., Garcia, J. A. Melton, J. A., Rogers, S. E., Shestopovlov, A. J., and Vicker, D. J., “Overflow Simulation Guidelines for Orion Launch Abort Vehicle Aerodynamic Analyses,” June 2011, AIAA Paper 2011-3163.
- [9] Rogers, S. E. and Pulliam, T. H., “Computational Challenges in Simulating Powered Flight of the Orion Launch Abort Vehicle,” *29th AIAA Applied Aerodynamics Conference*, June 2011, AIAA Paper 2011-3339.
- [10] Childs, R. E., Garcia, J. A., Rogers, S. E., and Vicker, D. J., “Overflow Aerodynamic Simulation of the Orion Launch Abort Vehicle,” *JANNAF 33rd Exhaust Plume and Signatures Meeting*, December 2012.
- [11] Vicker, D. J., Childs, R. E., Rogers, S. E., McMullen, M. S., Garcia, J. A., and Greathouse, J. S., “Effects of the Orion Launch Abort Vehicle Plumes on Aerodynamics and Controllability,” *51st AIAA Aerospace Sciences Meeting*, January 2013, AIAA Paper 2013-0970.
- [12] Nichols, R. H., Tramel, R. W., and Buning, P. G., “Solver and Turbulence Model Upgrades to OVERFLOW2 for Unsteady and High-Speed Applications,” *36th AIAA Fluid Dynamics Conference*, June 2006, AIAA Paper 2006-2824.
- [13] Chan, W. M., “Developments in Strategies and Software Tools for Overset Structured Grid Generation and Connectivity,” *20th AIAA Computational Fluid Dynamics Conference*, June 2011, AIAA Paper 2011-3051.
- [14] Rogers, S. E., Roth, K., Nash, S. M., Baker, M. D., P., S. J., Whitlock, M., and Cao, H. V., “Advances in Overset CFD Processes Applied to Subsonic High-Lift Aircraft,” *18th AIAA Applied Aerodynamics Conference*, June 2000, AIAA Paper 2000-4216.
- [15] Spalart, P. R. and Allmaras, S. R., “A One-Equation Turbulence Model for Aerodynamic Flows,” *Recherche Aerospatiale*, Vol. 1, 1994, pp. 5.
- [16] Menter, F. R., “Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications,” *AIAA Journal*, Vol. 32, No. 8, 1994, pp. 1598.

- [17] Menter, F. R., Kuntz, M., and Langtry, R., “Ten Years of Industrial Experience with the SST Turbulence Model,” *Turbulence, Heat and Mass Transfer*, Vol. 4, 2003, pp. 625.
- [18] Pandya, S. and Chan, W. M., “Computation of Sectional Loads from Surface Triangulation and Flow Data,” *20th AIAA Computational Fluid Dynamics Conference*, June 2011, AIAA Paper 2011-3680.
- [19] Aftosmis, M., Berger, M., and Adomavicius, G., “A Parallel Multilevel Method for Adaptively Refined Cartesian Grids with Embedded Boundaries,” *38th Aerospace Sciences Meeting*, January 2000, AIAA Paper 2000-0808.
- [20] Nemec, M. and Aftosmis, M., “Adjoint Error Estimation and Adaptive Refinement for Embedded-Boundary Cartesian Meshes,” *18th AIAA Computational Fluid Dynamics Conference*, June 2007, AIAA Paper 2007-4187.
- [21] Klopfer, G. H., Onufer, J. T., Pandya, S. A., Chan, W. M., Kless, J. E., and Lee, H. C., “Analyses of the Ares I A106-Plus Launch Vehicle First Stage Separation,” Tech. Rep. FS-TR-00002, NASA, 2011.